

UNIVERSITY OF CALIFORNIA SAN DIEGO

Data-Driven Forecasting for Grid-Connected Solar Power Plants

A dissertation submitted in partial satisfaction of the

requirements for the degree

Doctor of Philosophy

in

Engineering Sciences (Mechanical Engineering)

by

David Philip Larson

Committee in charge:

Professor Carlos F. M. Coimbra, Chair

Professor Jorge Cortés

Professor Philip E. Gill

Professor Jan P. Kleissl

Professor David G. Victor

2019

Copyright

David Philip Larson, 2019

All rights reserved.

The dissertation of David Philip Larson is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2019

DEDICATION

To my family.

EPIGRAPH

*It is not knowledge, but the act of learning,
not possession but the act of getting there,
which grants the greatest enjoyment.*

—Carl Friedrich Gauss

TABLE OF CONTENTS

Signature Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	ix
List of Tables	xi
Acknowledgements	xii
Vita	xiii
Abstract of the Dissertation	xiv
1 Introduction	1
1.1 Background	2
1.2 Motivation	3
1.3 Notation	4
1.4 Organization of the Dissertation	5
2 Day-Ahead Solar Power Forecasting	6
2.1 Prior Work	6
2.2 Data	9
2.2.1 Ground Data	9
2.2.2 NWP Data	11
2.3 Forecasting Models	11

2.3.1	Persistence	11
2.3.2	Deterministic Cloud Cover to GHI Model	13
2.3.3	Effects of Spatial Averaging	13
2.3.4	GHI to PO Model	14
2.4	Results and Discussion	17
2.4.1	Error Metrics	17
2.4.2	GHI Forecast Performance	18
2.4.3	Power Output Forecast Performance	18
2.4.4	Interannual Performance Variability	20
2.5	Summary	26
2.6	Acknowledgements	28
3	Intra-Day Solar Power Forecasting	29
3.1	Prior Work	29
3.1.1	Contributions	30
3.2	Forecast Methodology	31
3.2.1	Ordinary Least-Squares (OLS)	32
3.2.2	Support Vector Regression (SVR)	32
3.2.3	Persistence	33
3.2.4	Smart Persistence	34
3.3	Data	34
3.3.1	Satellite Images	35
3.4	Results and Discussion	38
3.4.1	Effect of Image Size	38
3.4.2	Effect of Model Choice	41
3.4.3	Comparison to Prior Work	42
3.5	Summary	45

3.6	Acknowledgements	47
4	Estimation of Cloud Optical Properties	48
4.1	Data	50
4.1.1	Remote sensing	50
4.1.2	Ground telemetry	52
4.2	Methods	54
4.2.1	SCOPE method	56
4.2.2	Spectral radiative model	57
4.2.3	Model tuning and sensitivity analysis	58
4.3	Results and Discussion	61
4.3.1	Daytime clear periods	64
4.3.2	Daytime cloudy periods	65
4.3.3	Daytime all-sky periods	66
4.3.4	Nighttime all-sky periods	66
4.4	Summary	67
4.5	Acknowledgements	68
5	Conclusions	74
	Bibliography	76

LIST OF FIGURES

Figure 1.1:	An annotated satellite view of the 250 MWp California Valley Solar Ranch. . .	4
Figure 2.1:	Availability of ground measurements of GHI and PO for CCA and LCC	10
Figure 2.2:	Map of Southern California, with NAM grid points, CCA, and LCC denoted by markers.	12
Figure 2.3:	Effect of spatial averaging the NAM_{GHI} forecasts for CCA.	15
Figure 2.4:	GHI vs. PO for CCA over September 2013–November 2014, ignoring night values ($\theta_z > 85^\circ$). Each marker is a single data point, and the line is a least-squares fit between GHI and PO with a R^2 value of 0.97.	16
Figure 2.5:	Sample days where the NWP-based forecasts ($RDPS_{CC}$ and NAM_{GHI}^*) for CCA perform (a) worse than; (b) as good as; and (c) better than day-ahead persistence. 21	21
Figure 2.6:	Power forecast error distributions for CCA, LLC, and treating the two sites as one combined power plant.	24
Figure 2.7:	Day-ahead PO forecast errors for CCA and LCC_j grouped by hour of the day and by month of the year.	25
Figure 3.1:	A diagram of the overall forecast methodology.	31
Figure 3.2:	A sample visible channel image from GOES-West.	35
Figure 3.3:	The merged PO and satellite image data sets for the two power plants: (a) CCA and (b) LCC. Two years of data (2012–2013) are designated for training the forecast models (shown in gray), with two additional years (2014–2015) used for testing (shown in black).	36
Figure 3.4:	Distribution of PO in the training and testing data sets for CCA and LCC. . . .	37
Figure 3.5:	Effect of image size on intra-day forecast accuracy.	40
Figure 3.6:	Kernel density estimates (KDE) of the 1-hour ahead SVR forecasts for CCA. .	43
Figure 3.7:	Same as Figure 3.6, but for the LCC testing set.	43

Figure 3.8:	Forecast error on the testing set, grouped by horizon.	46
Figure 4.1:	Illustration of the GOES-R Series satellite design.	51
Figure 4.2:	The spectral response function (SRF [-]) of the GOES-16 ABI longwave channels (7–16). For more details on the channels, see Table 4.1.	53
Figure 4.3:	The view of North America from the perspective of GOES-16.	55
Figure 4.4:	Illustration of the temperature profiles used in our method.	59
Figure 4.5:	Comparison of clear-sky OLR as a function of the number of atmospheric layers and spectral resolution of the radiative model.	60
Figure 4.6:	The upwelling spectral intensity as a function of T , ϕ , τ , and z_T	69
Figure 4.7:	Comparison of the normalized flux per channel.	70
Figure 4.8:	RMSE and MBE between the measured and modeled normalized flux per channel.	71
Figure 4.9:	RMSE and MBE between the measured and modeled DLW.	72
Figure 4.10:	The distribution of DLW absolute error.	73
Figure 4.11:	Same as Figure 4.9, but with the RMSE and MBE computed after removing samples with clouds that fully attenuate the incoming solar irradiance, using $k_t \leq 0.001$ as a threshold.	73

LIST OF TABLES

Table 2.1:	Summary of the current literature on day-ahead PO forecasting.	9
Table 2.2:	Day-ahead GHI forecast performance for CCA.	19
Table 2.3:	Day-ahead hourly PO forecasts for CCA and LCC.	20
Table 2.4:	Interannual variability of NWP-based PO forecasts for CCA and LCC.	23
Table 2.5:	Occurrence of over-predict errors.	26
Table 3.1:	Intra-day forecast performance on the testing set for CCA and LCC.	39
Table 3.2:	Intra-day forecast performance, broken down by site and season.	44
Table 4.1:	Summary of the ten longwave channels from the GOES-R Series ABI.	52
Table 4.2:	Summary of the seven SURFRAD stations.	54
Table 4.3:	Convergence of the clear-sky broadband OLR as the number of layers increases.	61

ACKNOWLEDGEMENTS

I would like to thank my advisor, Prof. Carlos F. M. Coimbra, for his support and guidance. I am forever grateful to the immeasurable positive impact he has had on my life.

I would also like to thank my committee members: Prof. Kleissl, Prof. Cortés, Prof. Gill, and Prof. Victor. I could not ask for a better group of faculty to oversee my doctoral work.

Last, but not least, I would like to thank my family, and the many friends and colleagues who have supported me throughout my time at UC San Diego.

Chapter 2 is taken, in part, from: D. P. Larson, L. Nonnenmacher and C. F. M. Coimbra, “Day-Ahead Forecasting of Solar Power Output from Photovoltaic Plants in the American Southwest”, *Renewable Energy* (2016), 91, pp. 11–20. The dissertation author was the primary investigator and author of this manuscript.

Chapter 3 is taken, in part, from: D. P. Larson and C. F. M. Coimbra “Direct Power Output Forecasts from Remote Sensing Image Processing”, *Journal of Solar Energy Engineering* (2018), 104 (2), pp. 021011-1–021011-08. The dissertation author was the primary investigator and author of this manuscript.

Chapter 4 is taken, in part, from: D. P. Larson, M. Li and C. F. M. Coimbra “Direct Estimation of Cloud Optical Properties and Downwelling Longwave Radiation from GOES-R Spectral Imagery” (in preparation). The dissertation author is the primary investigator and author of this manuscript.

VITA

- 2012 B. S. in Mechanical Engineering, University of California Merced
- 2014 M. S. in Engineering Sciences (Mechanical Engineering), University of California San Diego
- 2018 Ph. D. in Engineering Sciences (Mechanical Engineering), University of California San Diego

PUBLICATIONS

D. P. Larson, L. Nonnenmacher and C. F. M. Coimbra, “Day-Ahead Forecasting of Solar Power Output from Photovoltaic Plants in the American Southwest”, *Renewable Energy* (2016), 91, pp. 11–20. DOI: 10.1016/j.renene.2016.01.039

D. P. Larson and C. F. M. Coimbra, “Direct Power Output Forecasts from Remote Sensing Image Processing”, *Journal of Solar Energy Engineering* (2018), 104 (2), pp. 021011-1–021011-08. DOI: 10.1115/1.4038983

H. T. C. Pedro, D. P. Larson and C. F. M. Coimbra, “A Comprehensive Dataset for the Accelerated Development and Benchmarking of Solar Forecasting Methods”, *JRSE* (under review).

D. P. Larson, M. Li and C. F. M. Coimbra, “Direct Spectral Estimation of Cloud Optical Properties from GOES-R Imagery”, *JRSE* (in preparation).

ABSTRACT OF THE DISSERTATION

Data-Driven Forecasting for Grid-Connected Solar Power Plants

by

David Philip Larson

Doctor of Philosophy in Engineering Sciences (Mechanical Engineering)

University of California San Diego, 2019

Professor Carlos F. M. Coimbra, Chair

Forecasting is one of the enabling technologies for the integration of weather-dependent renewable resources (e.g., solar and wind) into the electric grid. Accurate forecasts can reduce operational costs associated with intra-day variability, reduce imbalance charges incurred by plant operators due to inaccurate energy bids, decrease utility costs associated with day-ahead scheduling (thereby reducing overall O&M costs), as well as assist grid operators with balancing energy demand schedules. As the market penetration of solar-based power generation continues to grow, accurate and reliable forecasting techniques become increasingly more important. In this work, two key areas

of solar forecasting are advanced. First, we develop intra-day (>1 -hour) and day-ahead (>24 -hour) forecasting methods to directly predict the generation of operational solar power plants, without the need for intermediate solar irradiance forecasts and resource-to-power modeling. Here we take a data-driven approach, leveraging Machine Learning (ML) techniques and publicly available, spatially resolved meteorological and remote sensing datasets. The proposed methods are analyzed and validated using two grid-connected 1 MW photovoltaic (PV) power plants in California. Second, we develop a method to directly and efficiently estimate cloud optical properties from longwave remote sensing data. The output of solar-based power generation systems is strongly dependent on cloud cover and optical depth, but in most solar forecasting methodologies cloud optical properties are over-simplified due to a lack of real-time, accurate estimates. The proposed estimation method builds upon a two-stream, spectrally resolved infrared radiation model coupled with high-resolution (5-minute, 2 km) spectral satellite imagery. We show that the proposed method can provide real-time, accurate estimates of cloud optical depth (COD) and cloud top height for all-sky (clear or cloudy) conditions during both daytime and nighttime.

Chapter 1

Introduction

There are a range of technological, societal and policy challenges associated with the world's transition to a low-carbon electric grid. A key technological challenge is the integration of intermittent renewable resources into the electric grid. While not a complete solution, forecasting provides an avenue towards enabling increasing penetration of renewable energy, particularly for non-dispatchable resources such as solar and wind. Accurate forecasts can reduce operational costs associated with intra-day variability, reduce imbalance charges incurred by plant operators due to inaccurate energy bids, decrease utility costs associated with day-ahead scheduling (thereby reducing overall O&M costs), and assist grid operators with balancing energy demand schedules. The value and importance of forecasting will continue increase as solar and wind power plants continue to grow, both in terms of the size of individual plants and the percentage of total generation capacity. In this work, we present contributions in two key areas related to solar forecasting: (1) data-driven methods to directly forecast the generation of solar-based power plants and (2) a method to directly estimate cloud optical properties in real-time, which can be used as inputs for modeling and forecasting applications.

1.1 Background

Here we provide a brief review of the key concepts related to solar forecasting. For a comprehensive review, readers are referred to [IPC13, LPKT14, AOU⁺16] and the references therein.

The goal of solar power forecasting is to predict the future output of a solar-based power generation system, e.g., a grid-connected photovoltaic (PV) power plant. In general, forecasts are based on current and past output of the system, i.e., endogenous information. However, the inclusion of exogenous variables, e.g., sensors or sky images, can provide significant improvements in forecasting performance by providing knowledge of factors that influence the power output, e.g., clouds. In addition, while it may be possible to describe the dynamics of portions of the system as differential equations, e.g., inverters, the chaotic and complex dynamics of weather phenomenon preclude this approach for solar forecasting. Instead, the majority of state-of-the-art forecasting methods tend to be based on data-driven, statistical methods.

Many solar forecasting methodologies have been developed for horizons ranging from minutes to hours to days ahead. Intra-hour forecasting ($t < 1$ h) is mainly based on sky imaging techniques and time-series models, while satellite-based forecasts achieve usable results for time intra-hour (1–6-hours) horizons [PKS⁺10, NC14, ZPC15]. Usually, horizons larger than 6 h require numerical weather prediction models (NWP) to generate accurate results, although there are exceptions such as [WMSZ12]. Recent advances in solar forecasting have mainly covered intra-day resource forecasts, driven by advances in sky image techniques and time-series modeling [CUL⁺11, MC13, GUY⁺14, CPN⁺14, CLPC15, BRAH14, PC15, CYL14, CA08, MZP⁺10].

While intra-day forecasts are important for grid stability, day-ahead forecasts are critical for market participation and unit commitment. Current market regulations in many world regions require day-ahead forecasts [KSHM13, LMK14, NKC15] and there are financial incentives to produce accurate forecasts [LMK14]. Besides market participation, day-ahead forecasts can also be useful for energy storage dispatch [HKNF14, DFM15].

To date, the majority of solar forecasting research has focused on the solar resource rather than power output (PO). In particular, the literature is biased towards global horizontal irradiance (GHI) forecasts, which are important for non-tracking photovoltaic (PV) power systems. There are studies on predicting direct normal irradiance (DNI) [NKC14, CPC13, CLPC15] as well as the PO of operational solar power plants [ZMAP14a, ZMAP14b, PGK13, LNC16] for intra-hour, intra-day and day-ahead horizons, but the literature is sparse compared to GHI.

1.2 Motivation

Increasing demand for renewable energy has led to rapid growth of large-scale ($\gg 1$ MWp) solar power plants, particularly in the American Southwest. In addition to their nameplate capacities, these power plants can also be considered large-scale due to their land area requirements. The California Valley Solar Ranch (CVSR) is a 250 MWp photovoltaic (PV) power plant in San Luis Obispo County and covers over 17 km² (see Figure 1.1). In the same region, the Topaz Solar Farm produces 550 MWp while using 25 km². To meet renewable energy integration goals, *e.g.*, California's 50% renewable by 2030, will require construction of additional large-scale power plants, of these sizes or larger.

The large-scale of these power plants result in unique challenges related to power forecasting. Specifically, the power plants cover areas large enough that the effects of solar microclimates can be seen in the power production. Unfortunately, cost and availability of instrumentation options prevents sufficient ground telemetry coverage at most power plants. To improve the modeling and forecasting performance of these large power plants requires higher spatial-resolution data, *e.g.*, sensor networks or remote sensing from next-generation satellites.

Another challenge for these large-scale plants relates to the need for real-time forecasting solutions. Deploying forecasts in real-time requires minimizing latency between all steps of the forecasting process, from data collection to communication to computation. In addition to technological sources of latency, *e.g.*, communication protocols, security requirements can inhibit



Figure 1.1: An annotated satellite view of the 250 MW_p California Valley Solar Ranch (CVSR). The sites covers approximately 17 km², although the area covered by solar panels is less than half of that size (≈ 8 km²).

or prevent implementation of real-time forecast methodologies. For example, power plants may require two-factor authentication for accessing power production data externally, thereby preventing automation of data collection needed for real-time forecast generation. These challenges should be considered when developing and evaluating forecast methods.

1.3 Notation

Each chapter of the dissertation discusses the notation relevant to the current chapter. However, here we make note of the general notation adopted for this dissertation. We adopt the notation commonly employed in the field of convex optimization, with a few minor deviations. Let \mathbf{R} , \mathbf{R}_+ and \mathbf{R}_{++} denote the set of real, nonnegative real, and positive real numbers. Unless noted, we use capital letters for matrices, *e.g.*, $A \in \mathbf{R}^{m \times n}$ is a $m \times n$ matrix. Lower case letters are vectors *e.g.*, $x \in \mathbf{R}^n$ is a n -length vector. Scalars and vectors are differentiated based on context. Vector norms are

written as $\|\cdot\|_p$, *e.g.*, $\|\cdot\|_2$ is the $\ell - 2$ norm. In the context of forecasting, true values are denoted without modification, while predicted values have hats, *e.g.*, y is the true value and \hat{y} is the forecasted value.

1.4 Organization of the Dissertation

Chapters 2 and 3 present work on data-driven solar power forecasting methods. Chapter 2 focuses on the day-ahead (>24 -hour) forecast horizon, while Chapter 3 targets the intra-day horizon (1–6-hours ahead). For both horizons, the methods are validated using two operational, 1 MW photovoltaic power plants in the American Southwest. Chapter 4 presents work on directly estimating cloud optical properties from satellite imagery.

Chapter 2

Day-Ahead Solar Power Forecasting

Day-ahead solar forecasts are critical for market participation and unit commitment. Current market regulations in many world regions require day-ahead forecasts [KSHM13, LMK14, NKC15] and there are financial incentives to produce accurate forecasts [LMK14]. Besides market participation, day-ahead forecasts can also be useful for energy storage dispatch [HKNF14, DFM15]. In this work, we focus on day-ahead forecasting of power output from photovoltaic (PV) power plants in the American Southwest.

2.1 Prior Work

Most NWP models generate global horizontal irradiance (GHI) at the ground level as one of the outputs, with some newer generation models now including direct normal irradiance (DNI). Previous studies evaluated the accuracy of this GHI variable and suggested ways to improve it, which have mostly focused on spatial averaging and bias removal [MK11, VMPN12, LFRAPV⁺12, MCK13, PLP⁺13, PGK13, OSF⁺13]. [LHHB09], [MK11] and [PGK13] showed that forecast errors for all sky conditions can be reduced by averaging the GHI forecasts from all NWP grid points within a set distance of the target site. In addition, [MK11] showed that the forecast performance could be further improved through bias removal using a polynomial regression based on the solar

zenith angle and clear sky index. [PGK13] showed that a similar improvement can be achieved through Kalman filtering.

In contrast to our knowledge of the performance of day-ahead GHI predictions, the application of these forecasts directly to prediction of day-ahead power output (PO) of PV plants is poorly understood. There are currently less than a dozen published studies that cover the subject of day-ahead PV PO forecasts, the majority of which were published in the last 5 years. We summarize these papers below and in Table 2.1. It is important to note that this lack of knowledge is partly due to the difficulty of obtaining data from operational PV plants, due to security restrictions and lack of data infrastructure. However, data access should improve in the coming years due to energy policies that require PO forecasting and therefore necessitate the collection of PO data.

[BMN09] applied the Danish Meteorological Institute's High Resolution Limited Area Model (HIRLAM) to forecast PO of 21 PV systems in Denmark. The PV systems had a total rated power capacity of 1–4 kWp each and one year of data was used to evaluate the forecasts. HIRLAM GHI forecasts were used as the main input to autoregressive-based PO models. [MBCCM⁺14] also used HIRLAM to forecast PO of a 6.2 kWp test site in Spain.

[LHHB09] and [LSH⁺11] used the European Center for Medium-Range Weather Forecasts Model (ECMWF) to generate regional, PO forecasts in Germany. [LHHB09] only evaluated two months (April and July 2006) of forecasts for 11 PV systems in Southern Germany, while [LSH⁺11] tested on 10 months (July 2009–April 2010) of PO data from approximately 400 representative PV systems.

[MP10] used an artificial neural network (ANN) to predict GHI in Trieste, Italy. The predicted GHI was then mapped directly to PO using the efficiency data of the studied 20 kWp PV system. Unfortunately, PO forecast results were only reported for four consecutive clear sky days in May 2009. Similarly, [LSS14] evaluated an autoregressive–moving-average model with exogenous inputs (ARMAX) forecast model that did not use NWP data as input. However, [LSS14] only forecasted the day-ahead mean daily PO, rather than day-ahead hourly values, for a 2.1 kWp PV system in China.

[PGK13] forecasted PO of three small PV systems (6.72, 19.8 and 45.6 kWp) in mainland Canada, using the Canadian Meteorological Centre's Global Environmental Multiscale Model (GEM). The GEM's GHI forecasts were validated against ground measurements from the SURFRAD Network in the United States. Spatial averaging of 300km to 600km and bias removal via Kalman filtering were used to improve the GEM forecast performance. Reported RMSE values were in the range of 6.4% to 9.2% of the rated power of the PV systems for a 1 year testing set (April 2008–March 2009).

[FOO⁺14a] and [FOO⁺14b] generated regional forecasts for Japan using Support Vector Regression (SVR) together with inputs from the Japan Meteorological Agency's Grid Point Value-Mesoscale Model (GPV-MSM). [FOO⁺14a] used PO data from approximately 450 PV plants from four regions (Kanto, Chubu, Kansai, Kyushu) that had a net rated power of approximately 15 MWp while [FOO⁺14b] used data from 273 PV plants spread over two regions (Kanto and Chubu), with an approximate total power of 8 MWp.

[ZMAP14a] and [ZMAP14b] forecasted PO of 28 PV plants in mainland France using Meteo France's Action de Recherche Petite Echelle Grande Echelle (ARPEGE) model. [ZMAP14a] presented a deterministic forecast that achieved a RMSE of 8–12% of the plant capacity over two years of testing data (2009–2010) while [ZMAP14b] presented a probabilistic forecast. Both methods used 31 variables from the ARPEGE, including GHI as well as environmental conditions, e.g., temperature, humidity and precipitation. Unfortunately, no actual values were reported for the PV plant power ratings or other technical details, which limits analysis into the applicability of the results to other PV plants and regions.

Most recently, [APN15] forecasted the PO of five vertical-axis tracking PV plants in Spain, using a nonparametric model based on the Weather Research and Forecasting Model (WRF). The PV plants ranged in size from 775 to 2000 kWp and the forecasts were evaluated with data from 2009–2010. Quantile Regression Forests, a variation of random forests, was used to generate the PO forecasts, with WRF variables such as GHI, temperature and cloud cover as the inputs.

Although these studies all presented day-ahead PO forecasting for PV plants, further research

Table 2.1: Summary of the current literature on day-ahead PO forecasting. “Type” refers to whether the forecast was for a single PV plant site or for an entire region while “Testing data” is how much data was used in the forecast model evaluation.

	Region	NWP model	Type	PV plants		
				No. of sites	Total capacity	Testing data
[BMN09]	Denmark	HIRLAM	point	21	~100 kWp	1 year
[MBCCM ⁺ 14]	Spain	HIRLAM	point	1	6 kWp	10 months
[LHHB09]	Germany	ECMWF	regional	11	unknown	2 months
[LSH ⁺ 11]	Germany	ECMWF	regional	383	~100 MWp	10 months
[FOO ⁺ 14a]	Japan	GPV-MSM	regional	454	~10 MWp	1 year
[FOO ⁺ 14b]	Japan	GPV-MSM	regional	273	~10 MWp	1 year
[PGK13]	Canada	GEM	point	3	~100 kWp	2 years
[ZMAP14a]	France	ARPEGE	point	28	unknown	2 years
[ZMAP14b]	France	ARPEGE	point	28	unknown	2 years
[APN15]	Spain	WRF	point	5	~10 MWp	2 years
[MP10]	Italy	N/A	point	1	20 kWp	4 days
[LSS14]	China	N/A	point	1	2 kWp	6 months

is still required, especially for sites in the United States. In this study we seek to provide the following contributions: (1) Introduction and evaluation of a PV PO forecast model for the American Southwest, a region with both high solar energy generation potential and a favorable political environment for solar, especially in California. (2) Investigation of the interannual forecast performance variability and (3) the occurrence of severe forecasting errors, as they relate to large-scale renewable energy integration. (4) Spatial smoothing through pairing of PV plants in proximity. To achieve these goals, we systematically test approaches and inputs to generate PO forecasts based on two operational PV plants in Southern California.

2.2 Data

2.2.1 Ground Data

Two sites are used to evaluate day-ahead forecasts of GHI and PO: Canyon Crest Academy (CCA) and La Costa Canyon High School (LCC). CCA (32.959°, -117.190°) and LCC (33.074°, -117.230°) are both located in San Diego County, USA, and each feature 1 MW peak (MWp) of non-tracking photovoltaic (PV) panels. The PV panels for both sites were installed in 2011 as part

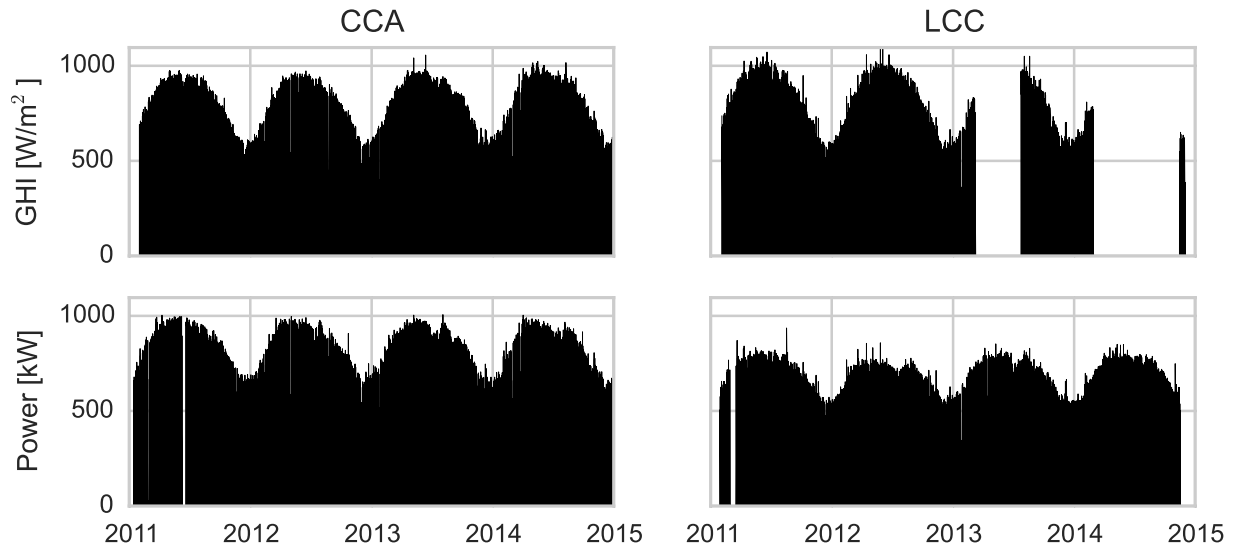


Figure 2.1: Availability of ground measurements of GHI and PO for CCA (left column) and LCC (right column). The missing GHI data from LCC is due to a malfunctioning pyranometer.

of a power purchase agreement (PPA) between the San Dieguito Union High School District and Chevron Energy Solutions.

Each site has two inverters, with the PV panels mounted above the parking lots at a 5° incline. At CCA, the panels are split over two parking lots, approximately 200m apart and each with 500 kWp. Due to the site's characteristics, the north parking lot panels face ($\sim 20^\circ$) southwest, while the south parking lot panels face directly south. Meanwhile, the panels at LCC cover a single parking lot and face $\sim 30^\circ$ southwest.

Total PO from the PV panels at each site is available as 15-minute averages. Additionally, 15-minute averaged GHI data is available from co-located pyranometers. Figure 2.1 illustrates the data availability for both sites. For this study, the GHI and PO data are backwards-averaged to hourly values and night values are removed, where we define night as zenith angles (θ_z) greater than 85° . Note that the maximum observed PO of CCA matches the nameplate capacity (1 MWp), while the maximum observed PO of LCC peaks at ≈ 0.8 MW, despite both sites using identical PV technologies. This difference is due to the greater ($\sim 30^\circ$) southwest orientation of the PV panels at LCC.

2.2.2 NWP Data

We use forecasted variables from two publicly-available NWP models: the North American Mesoscale Forecast System (NAM) and the Regional Deterministic Prediction System (RDPS).

The NAM is a NWP model provided by the National Oceanic and Atmospheric Administration (NOAA) on a $\approx 12 \text{ km} \times 12 \text{ km}$ spatial grid that covers the continental United States. Forecasts are generated four times daily at 00Z, 06Z, 12Z and 18Z, with hourly temporal resolution for 1–36 h horizons and 3 h resolution for 39–84 h horizons. Downward shortwave radiative flux (DSWRF) [W/m^2] at the surface, a synonym for GHI, is forecasted using the Geophysical Fluid Dynamics Laboratory Shortwave radiative transfer model (GFDL-SW) [LH74]. The NAM also forecasts total cloud cover (TCDC) [%], where the entire atmosphere is treated as a single layer. In this study, NAM forecasts generated at 00Z were downloaded from the NOAA servers and degribbed for September 2013 to November 2014, for all nodes within 200km of CCA and LCC. Figure 2.2 shows the distribution of NAM forecast nodes over Southern California.

The Canadian Meteorological Centre generates the RDPS model on a $10 \text{ km} \times 10 \text{ km}$ spatial grid covering North America. As with the NAM, RDPS forecasts total cloud cover (TCDC) as a percentage for each grid element, not distinguishing between the layers of the atmosphere. The operational RDPS model generates forecasts daily at 00Z, 06Z, 12Z and 18Z with hourly temporal resolution. Only the 00Z forecasts are used in this study, to ensure fair comparison with the NAM forecasts. RDPS GHI forecasts were not available.

2.3 Forecasting Models

2.3.1 Persistence

We include a day-ahead persistence forecast of both GHI and PO as a baseline. The persistence forecast assumes that the current conditions will repeat for the specified time horizon. In

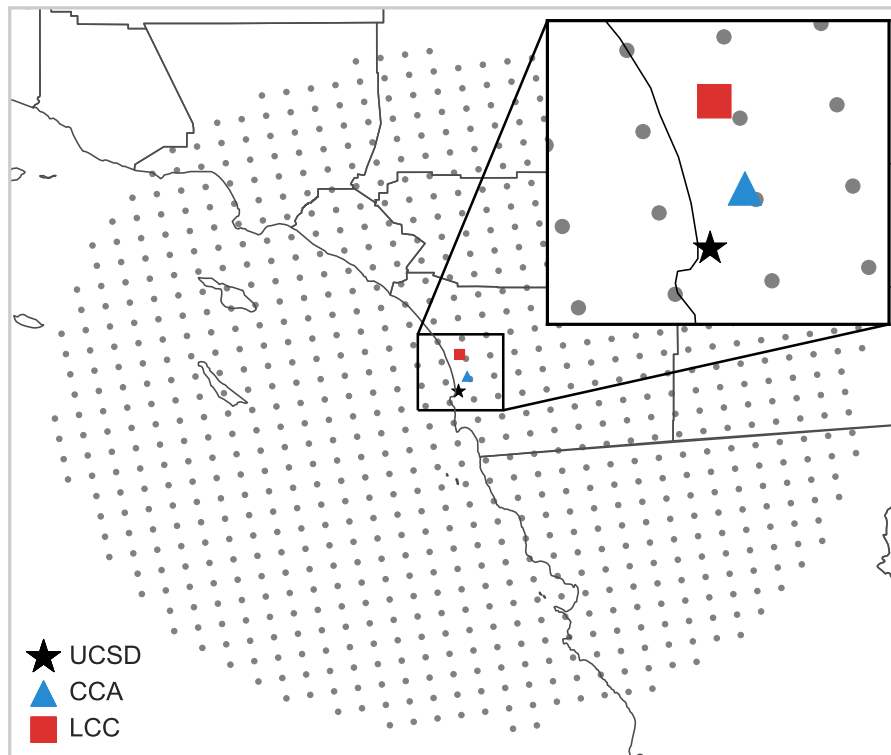


Figure 2.2: Map of Southern California, with the NAM grid points marked with gray dots. UC San Diego (UCSD; star), Canyon Crest Academy (CCA; triangle) and La Costa Canyon (LCC; square) are also marked on the map. The area to the lower left is the Pacific Ocean, with Mexico shown in the bottom right.

other words:

$$\hat{y}_{t+\tau} = y_t, \quad (2.1)$$

where $\hat{y}_{t+\tau}$ is the forecasted value of some variable y , at a time horizon of τ from some initial time t . For day-ahead persistence, $\tau = 24h$.

2.3.2 Deterministic Cloud Cover to GHI Model

We use a modified version of the deterministic model from [NKC15], [MC11] and [PMW⁺07] to derive GHI from cloud cover:

$$\text{GHI} = \text{GHI}_{\text{CS}}[0.35 + 0.65(1 - \text{CC})], \quad (2.2)$$

where CC is the TCDC (0=clear, 1=overcast) from the NWP model and GHI_{CS} is the clear sky GHI from the airmass independent model described in [IP02], [Ine02], and [Gue12].

2.3.3 Effects of Spatial Averaging

Previous studies have shown that NAM GHI forecast error can be reduced by spatially averaging the forecasts from all nodes within a set distance of the target site [MK11, PGK13]. The distance (d) in km between a site and a NAM node can be calculated using the spherical law of cosines:

$$d = \arccos(\sin(\phi_1) \sin(\phi_2) + \cos(\phi_1) \cos(\phi_2) \cos(\lambda_2 - \lambda_1)) R, \quad (2.3)$$

where (ϕ_i, λ_i) is the latitude-longitude for a point i and R is the radius of the Earth (6371 km). For both CCA and LCC, there are approximately 50 NAM nodes within 50km, 200 nodes within 100km, and 800 nodes within 200km.

Figure 2.3 shows the effect spatial averaging of the NAM_{GHI} GHI forecasts for CCA compared to using only the NAM node closest to the site, i.e., the naive node choice. The NAM forecasts from the available data set (September 2013–November 2014) are grouped into bins based on the clear sky index (k_t),

$$k_t = \frac{\text{GHI}}{\text{GHI}_{\text{CS}}}, \quad (2.4)$$

which provides a measure of the sky conditions during the forecasted time period ($k_t \rightarrow 0$: overcast, $k_t \rightarrow 1$: clear). Then, the root mean square error (RMSE) is calculated between the forecasts in each k_t bin and the ground truth, and averaged to produce one RMSE value per k_t bin. The result is that spatial averaging reduces the RMSE of the forecasts for all sky conditions, as compared to the naive node choice. This is consistent with results presented in previous work, e.g., [MK11] and [PGK13].

Hereafter, we will report results for the naive NAM forecasts (NAM_{GHI} , NAM_{CC}) as well as the 100km spatially averaged versions ($\text{NAM}_{\text{GHI}}^*$, NAM_{CC}^*). 100km is chosen over 200km as there is a negligible difference in error between the two. Also, the 100km case requires approximately 25% the NAM grid points as the 200km case, and therefore has lower storage and computational requirements.

2.3.4 GHI to PO Model

As the PV panels at both sites are non-tracking, we assume that there exists a linear function $f(\cdot)$ that can map GHI to PO. Figure 2.4 illustrates the relationship between GHI and PO for CCA over a full year (September 2013–November 2014), with a reference linear fit included ($R^2 = 0.97$). The high R^2 value (> 0.95) indicates that a linear mapping between GHI and PO is not an unreasonable assumption.




Figure 2.3: Effect of spatial averaging the NAM_{GHI} GHI forecasts for CCA September 2013–November 2014. Naive refers to the NAM node physically closest to the site, while 50km, 100km, and 200km are the unweighted average forecasts of all nodes within each search radius. Each marker represents the average RMSE of all forecasts generated for each k_t bin, i.e., each group of sky conditions.




Figure 2.4: GHI vs. PO for CCA over September 2013–November 2014, ignoring night values ($\theta_z > 85^\circ$). Each marker is a single data point, and the line is a least-squares fit between GHI and PO with a R^2 value of 0.97.

Formally, we define a linear function $f(x; w)$ as:

$$f(x; w) = \sum_{j=0}^n w_j x_j = w^T x, \quad (2.5)$$

where w is the weight vector, x is the input vector, $x_0 = 1$ by convention, and n is the number of model parameters. Then we use ordinary least-squares to find the optimal weights w^* , i.e., the weights that minimize the error between the predictor $f(x; w)$ and the true value y .

More specifically, for a given set of m example inputs $\{x^{(i)}\}_{i=1}^m$ and outputs $\{y^{(i)}\}_{i=1}^m$, the least-squares costs function is defined as:

$$J(w) = \sum_{i=1}^m \left(f(x^{(i)}; w) - y^{(i)} \right)^2 = \|Xw - y\|_2^2. \quad (2.6)$$

The optimal weights w^* are then the w that minimizes the cost function:

$$w^* = \arg \min_w \|Xw - y\|_2^2, \quad (2.7)$$

where $X \in \mathbf{R}^{m \times n}$, $w \in \mathbf{R}^n$, and $y \in \mathbf{R}^m$.

For the purposes of this study, we use the predicted GHI ($\widehat{\text{GHI}}$) from one of the NWP models and zenith angle (θ_z) as the model inputs. I.e. $n = 2$, $x_1 = \widehat{\text{GHI}}$ and $x_2 = \theta_z$, where θ_z is included so that temporal information is encoded into the model. Readers are referred to [SDBF⁺15] and [DLM15] for more detailed studies into modeling of PO of PV power plants.

2.4 Results and Discussion

2.4.1 Error Metrics

We use standard error metrics to evaluate the performance of both the GHI and PO forecasts: Mean Absolute Error (MAE), Mean Bias Error (MBE), Root Mean Square Error (RMSE), and skill (s) [MPC13]:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |e_i|, \quad (2.8)$$

$$\text{MBE} = \frac{1}{N} \sum_{i=1}^N e_i, \quad (2.9)$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N e_i^2}, \quad (2.10)$$

$$s = 1 - \frac{\text{RMSE}}{\text{RMSE}_p}, \quad (2.11)$$

where N is the number of forecasts, RMSE_p is the RMSE of the persistence forecast, and e is the error between the true value (y) and forecasted value (\hat{y}):

$$e_i = y_i - \hat{y}_i. \quad (2.12)$$

Additionally, for ease of comparison to previously published results, we also include normalized RMSE (nRMSE):

$$\text{nRMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N \left(\frac{e_i}{\bar{y}} \right)^2}, \quad (2.13)$$

where \bar{y} is the mean true value over the considered period, e.g., September 2013–November 2014.

2.4.2 GHI Forecast Performance

Day-ahead GHI forecast performance and statistics for CCA for September 2013–November 2014 are shown in Table 2.2. As noted in Section 2.3.3, the spatially averaged forecasts ($\text{NAM}_{\text{GHI}}^*$ and NAM_{CC}^*) use all NAM nodes within 100km of the site. The best GHI forecast is generated with RPDS, which reduces the RMSE 24% as compared to the persistence model. NAM_{CC} performs the worst, reducing RMSE only by 3% compared to persistence. However, the results presented are for GHI forecasts that have not undergone any post-processing or model output statistics (MOS), which have been shown to improve forecast performance (see [MK11, MCK13]).

Without post-processing or MOS, the GHI forecasts underperform as compared to prior literature. Two of the most extensive studies on day-ahead GHI forecasting are [PGK13] and [PLP⁺13]. In [PGK13], the GEM NWP model is evaluated using seven SURFRAD Network sites in the US and achieves RMSE values in the range of $\approx 80 - 120 \text{ Wm}^{-2}$. [PLP⁺13] expands on [PGK13] by providing results for a variety of common NWP models, including GEM, ECMWF and WRF. RMSE values reported in [PLP⁺13] are in the range of 70–200, 90–130, 80–115, and 95–135 Wm^{-2} for sites in the US, Central Europe, Spain and Canada respectively. Unfortunately, neither study directly evaluated RDPS or NAM.

2.4.3 Power Output Forecast Performance

PO is forecasted for both CCA and LCC using the methods described in Section 2.3. For the 18 months of overlapping data between the ground measurements and NWP models, odd months are

Table 2.2: Day-ahead GHI forecast performance for CCA, from September 2013 to November 2014. Night values have been removed (night $\equiv \theta_z \geq 85^\circ$). No post-processing or model output statistics (MOS) have been applied to the forecasts. The errors are relative to the mean GHI of the considered period (515 Wm^{-2}).

Method	MAE [%]	MBE [%]	RMSE [%]	s [-]
Persistence	24.9	7.7	36.2	—
RDPS _{CC}	20.3	2.3	27.5	0.24
NAM _{GHI}	21.9	-11.6	31.8	0.12
NAM [*] _{GHI}	20.5	-7.6	28.3	0.22
NAM _{CC}	26.4	14.1	35.3	0.03
NAM [*] _{CC}	26.2	16.2	32.7	0.10

used for the training set and even months for the testing set. Splitting the sets in this way, rather than in sequential partitions, ensures that the models are trained on data from all four seasons. Besides this monthly training, the effects of training on a full year of data and testing on another full year are discussed in Section 2.4.4.

Figure 2.5 shows three exemplary days of forecasted and measured PO. The persistence model performs well on the first (clear) day, with no visible difference to the measured PO, while both NWP-based forecasts under-predict the PO. On the second exemplary day, the NWP-based models perform better than on the first day, but do not beat the persistence model. Since the persistence model is just a repetition of the previous day, persistence works well if the weather conditions between days do not change. Hence, persistence also outperforms the NWP models because it inherits a better representation of the temperature effects on the panels that are not an input to the GHI to PO model. However, the strength of the NWP models is to predict cloud cover and irradiance attenuation, and are hence performing better when conditions change between days. The third exemplary day shows this effect, where the persistence model under-predicts since the previous day (not shown) was overcast. These results are expected and consistent with previous findings on the performance of the persistence model [NC14, LFRAPV⁺12].

Table 2.3 summarizes the PO forecast results and shows that all three NWP-based forecast models perform comparably for both sites ($s = 19 - 25\%$) and overall outperform the persistence model. However, the persistence model has a lower absolute error than the NWP-based models for

Table 2.3: Day-ahead hourly PO forecasts for CCA and LCC, with MAE, MBE and RMSE reported as relative to the rated plant capacity of each site (1 MWp). The models were trained and tested using hourly data from September 2013–November 2014, with odd months used for training and even months for testing.

Site	Method	Training				Testing			
		MAE [%]	MBE [%]	RMSE [%]	s [-]	MAE [%]	MBE [%]	RMSE [%]	s [-]
CCA	Persistent	9.0	0.4	15.6	—	8.8	-0.5	14.7	—
	RDPS _{CC}	8.9	0.0	12.5	0.20	8.4	0.2	11.5	0.21
	NAM _{GHI}	9.6	-0.0	13.2	0.15	8.5	-0.3	11.3	0.23
	NAM* _{GHI}	9.3	0.0	12.7	0.18	8.4	-0.2	11.0	0.25
	NAM _{CC}	9.2	-0.0	12.8	0.18	8.8	0.1	11.6	0.21
	NAM* _{CC}	9.1	-0.0	12.5	0.20	8.5	0.4	11.2	0.23
LCC	Persistent	8.0	0.3	14.0	—	7.2	-0.3	12.0	—
	RDPS _{CC}	7.7	0.0	11.1	0.21	7.3	0.6	9.8	0.19
	NAM _{GHI}	8.1	0.0	11.6	0.17	7.3	0.2	9.5	0.21
	NAM* _{GHI}	8.0	0.0	11.2	0.20	7.2	0.3	9.3	0.23
	NAM _{CC}	8.1	0.0	11.4	0.18	7.7	0.4	9.8	0.18
	NAM* _{CC}	8.0	0.0	11.2	0.20	7.3	0.6	9.5	0.21

47% and 49% of the data points in the testing set for CCA and LCC respectively. This is likely a result of the temperate climate at both sites, i.e., that both sites have an abundance of periods where the weather conditions do not change and therefore the persistence model will excel.

Given the significant variations in skill for the GHI forecasts ($s = 10 - 24\%$), it might be surprising that all three PO forecasts perform almost equally well. However, the training of the GHI to PO model removes biases in the forecasts. This seems to be an advantage of our proposed model compared to [PGK13], where the accuracy of the PV output prediction was strongly impacted by the accuracy of the GHI input.

2.4.4 Interannual Performance Variability

While all previous work has only focused on bulk error statistics, we also investigate the interannual performance variability. This is important since meteorological studies have shown that there are relatively strong interannual variations in many weather phenomenon, e.g., DNI [GW11]. However, this variation is unknown for the evaluation of PV PO. Quantifying this variation will

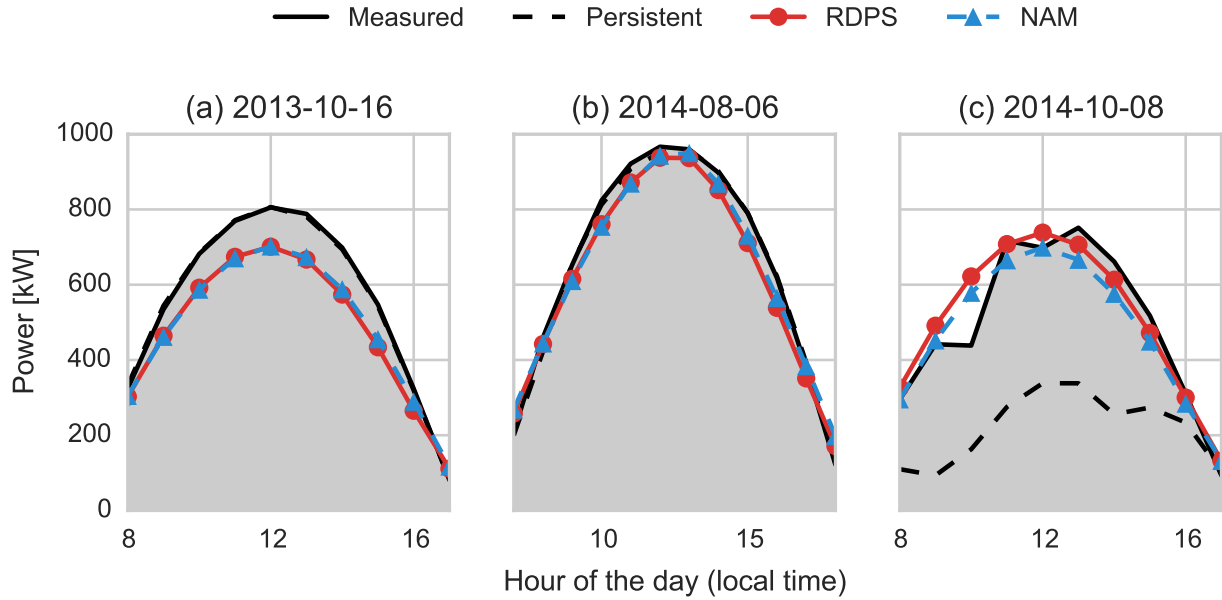


Figure 2.5: Sample days where the NWP-based forecasts (RDPS_{CC} and NAM_{GHI}^{*}) for CCA perform (a) worse than; (b) as good as; and (c) better than day-ahead persistence.

enable us to define what is a sufficient length of training and testing data sets for PO forecasts.

We compare the PO forecast performance per year for 2011–2014. RDPS_{CC} is used as the input source for the GHI to PO model and the model parameters are trained in the same way as the previous section. We use RDPS_{CC} rather than NAM_{GHI} and NAM_{CC} for two reasons. First, we have RDPS_{CC} forecasts for the entirety of 2011–2014, but NAM forecasts only for September 2013–November 2014. Second, as discussed in the previous section, the PO forecast performance is not strongly influenced by the input GHI forecast. Hence, we assume that the interannual RDPS_{CC} forecast results in this section are representative of all evaluated NWP-based models.

Table 2.4 shows the interannual performance variability of PO forecasts for CCA and LCC over 2011–2014. The PO model is trained separately on 2011, 2012, and 2013, and then tested on the subsequent years of data. For both CCA and LCC, the forecast performance on the test set depends more on the test set itself rather than the training set used. For example, the CCA forecast models trained on 2011, 2012 and 2013 achieved nRMSE values of 0.26, 0.30 and 0.28 respectively on the training set. But all three models achieved nRMSE of 0.24 on the 2014 testing set. A similar

trend is seen in the 2013 testing sets and the LCC forecasts.

Although there is no prior literature that reports day-ahead PO forecast performance in our studied region of the American Southwest, it is important to consider our results in the context of other work. As noted in Table 2.4, our forecast methodology achieves nRMSE values in the range of 0.24–0.33. [BMN09], [FOO⁺14b], and [APN15] reported nRMSE values in the range of 0.08–0.13, 0.26–0.35, and 0.10–0.30 respectively. Each study used at least 1 year of testing data, but the 21 PV systems in [BMN09] had nameplate capacities of only 1–4 kWp. [FOO⁺14b] and [APN15] evaluated forecasts with PV systems of total capacity ~ 10 MWp, and achieved nRMSE values closer in range to our results.

Error distributions are another important factor in evaluating forecast performance. Figure 2.6 shows the forecast error distributions of three training-testing data sets and three sites: CCA, LCC and a third fictional site denoted as “Combined”. The “Combined” site allows us to consider the impact of pairing the two PV plants to simulate a single, spatially distributed 2 MWp power plant. All three sites and training-testing sets show similar distributions, with peaks centered at less than zero error. 2012 and 2013 have a higher occurrence of over-prediction (negative) errors than 2014.

In addition to the annual statistics and error distributions shown in Table 2.4 and Figure 2.6, it is useful to evaluate the error distributions on shorter time scales. Figure 2.7 examines the error distribution of both sites broken down by the hour of the day and by the month of the year for the 2014 testing set. Violin plots are used instead of box plots to enable visualization of not only the error magnitudes, but also their distribution densities [HN98]. From the hour of day perspective, forecasts for both sites over-predict at sunrise and sunset, and under-predict during the majority of the day. On the monthly time scales, the forecasts tend to under-predict during the rainy season (December–March), when the coastal areas are more impacted by mobile mid-latitude weather systems and their associated cloud patterns than by the diurnal pattern of marine stratus coverage.

For system operators, over-prediction of power can be more harmful than under-prediction due to the technical complexity required for up-ramping of load following and backup units in the power system, which is greater than for curtailing excess output from intermittent producers.

Table 2.4: Interannual variability of NWP-based PO forecasts for CCA and LCC. RDPS_{CC} is used as the NWP input for the forecast, with night values removed ($\text{night} \equiv \theta_z > 85^\circ$). The GHI-to-PO model is trained on one year of data and then tested on a separate year. RMSE is reported as relative to the PV plant capacity (1 MWp for CCA and LCC) while nRMSE is calculated using Equation 2.13 and the mean PO ($\overline{\text{PO}}$) of the considered period. Higher forecast skills (s) indicate better performance.

Training					Testing				
Year	$\overline{\text{PO}}$ [kW]	RMSE [%]	s [-]	nRMSE [-]	Year	$\overline{\text{PO}}$ [kW]	RMSE [%]	s [-]	nRMSE [-]
CCA:									
2011	484	12.6	0.23	0.26	2012	466	14.0	0.17	0.30
2011	484	12.6	0.23	0.26	2013	476	13.7	0.15	0.29
2011	484	12.6	0.23	0.26	2014	481	11.7	0.23	0.24
2012	466	13.9	0.17	0.30	2013	476	13.6	0.16	0.29
2012	466	13.9	0.17	0.30	2014	481	11.7	0.22	0.24
2013	476	13.6	0.16	0.28	2014	481	11.7	0.23	0.24
LCC:									
2011	397	11.3	0.20	0.28	2012	374	12.4	0.13	0.33
2011	397	11.3	0.20	0.28	2013	392	11.7	0.14	0.30
2011	397	11.3	0.20	0.28	2014	410	10.2	0.20	0.25
2012	374	12.2	0.15	0.33	2013	392	11.8	0.13	0.30
2012	374	12.2	0.15	0.33	2014	410	10.5	0.17	0.26
2013	392	11.7	0.14	0.30	2014	410	10.3	0.20	0.25




Figure 2.6: Power forecast error distributions for CCA (top row), LCC (middle row) and treating the two sites as one combined power plant (bottom row). The columns are for three separate testing sets (2012, 2013, and 2014), with the PO forecast trained on the year prior to the testing set (e.g. 2012 was trained on 2011). As with Figure 2.7, the error is calculated using Equation 2.12 relative to the plant capacity, where negative errors are over-predictions (i.e. the forecasted PO was greater than the measured PO) while positive values are under-predictions.

Meanwhile, for power producers, curtailing solar PO is unfavorable as the power potential and profits of the plant are reduced. To better understand the occurrence of such errors, we analyze the severe over-prediction errors per year, both on hourly and daily timescales.

Table 2.5 summarizes the occurrence of over-prediction errors that exceed a range of thresholds, relative to the rated plant capacity. We define over-prediction errors greater than 20% as severe. Although both sites are approximately 7km from the Pacific Ocean and 13km from each other, LCC has a lower frequency of over-prediction events than CCA on both hourly and daily timescales. And

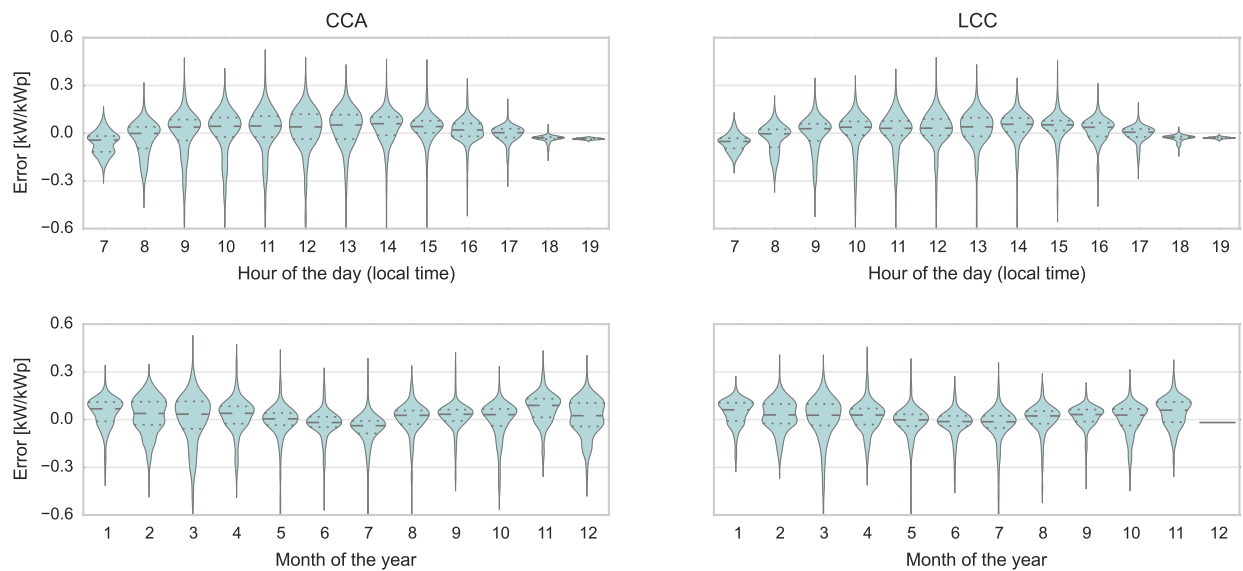


Figure 2.7: Day-ahead PO forecast errors for CCA (left column) and LCC (right column), grouped by hour of the day (top row) and by month of the year (bottom row). The results here are for the PO model trained on 2013 and then tested on 2014. Error is as defined in Equation 2.12 relative to the PV plant capacity. Therefore, negative error values are over-predictions by the PO model and positive values are under-predictions. The dashed lines show the 25th, 50th, and 75th percentiles, while the width represents the density, which is approximated using kernel density estimation (KDE).

Table 2.5: Occurrence of over-prediction errors that exceed 20%, 30%, 40%, and 50% of the plant capacity. The year denotes the year used for testing, with the previous year used for training. “Hourly” are the number of hourly data points, “Daily” are the number of days where the average of the hourly forecasting errors within each day exceed the threshold, and the values in parentheses are occurrences relative to the total number of data points [%]. 2012 and 2013 had approximately 4000 hourly and 365 daily data points per site, while 2014 had approximately 3600 hourly and 320 daily data points per site. “Combined” represents the scenario in which CCA and LCC are treated as a single, spatially distributed 2 MWp power plant.

Error	Hourly			Daily		
	CCA	LCC	Combined	CCA	LCC	Combined
2012:						
>20%	347 (8.6)	272 (6.7)	288 (7.1)	22 (6.0)	18 (4.9)	21 (5.7)
>30%	182 (4.5)	146 (3.6)	154 (3.8)	11 (3.0)	7 (1.9)	8 (2.2)
>40%	100 (2.5)	73 (1.8)	83 (2.1)	1 (0.3)	0 (0.0)	0 (0.0)
>50%	53 (1.3)	22 (0.5)	35 (0.9)	0 (0.0)	0 (0.0)	0 (0.0)
2013:						
>20%	302 (7.6)	208 (5.2)	228 (5.7)	17 (4.7)	7 (1.9)	8 (2.2)
>30%	143 (3.6)	85 (2.1)	100 (2.5)	4 (1.1)	0 (0.0)	2 (0.5)
>40%	66 (1.7)	28 (0.7)	39 (1.0)	0 (0.0)	0 (0.0)	0 (0.0)
>50%	24 (0.6)	6 (0.2)	14 (0.4)	0 (0.0)	0 (0.0)	0 (0.0)
2014:						
>20%	212 (5.9)	162 (4.5)	168 (4.7)	10 (3.1)	7 (2.2)	8 (2.5)
>30%	85 (2.4)	61 (1.7)	69 (1.9)	2 (0.6)	2 (0.6)	2 (0.6)
>40%	38 (1.1)	27 (0.8)	30 (0.8)	1 (0.3)	0 (0.0)	0 (0.0)
>50%	21 (0.6)	13 (0.4)	16 (0.4)	0 (0.0)	0 (0.0)	0 (0.0)

this trend is consistent for all three years of testing data (2012, 2013 and 2014).

The pairing of the two sites (“Combined”) has an impact on the occurrence of severe over-predictions. On the hourly timescale, the occurrence of severe over-prediction events for the “Combined” site is reduced for all thresholds by 0.2–1.9% compared to CCA. Similarly, daily over-prediction events are reduced by up to 2.5%. These results indicate that the pairing of PV plants can enable spatial smoothing of severe over-prediction events.

2.5 Summary

This study presented a methodology to generate day-ahead power output forecasts for two PV plants in the American Southwest. The forecasts are based on publicly available numerical

weather prediction models from the National Oceanic and Atmospheric Administration, and the Canadian Meteorological Centre. Four years of ground measurements (2011–2014) from two 1 MWp, non-tracking PV plants in San Diego County, USA were used in this study. Forecasts of the two sites achieved annual RMSE of 11.7–14.0% and 10.2–12.4% relative to the rated plant capacities (1 MWp), as well as annual forecast skills of 15–23% and 13–20%.

Our analysis leads to the following main conclusions:

1. Our proposed methodology can reduce the error in PV power output forecasts for day-ahead market participation in San Diego County, a key region for solar installations in the American Southwest.
2. Our methodology can be applied without any prior knowledge of the NWP model that produced the input GHI forecast.
3. The interannual forecast performance variability depends more heavily on the data set used in the performance evaluation rather than the data set used in training the forecast model.
4. Evaluation of severe over-prediction errors should be considered when evaluating day-ahead PO forecasts. Our methodology results in severe ($> 20\%$ of the nameplate capacity) hourly over-prediction errors less than 9% of the time over three years for both sites. However, the occurrence of such severe over-predictions may be larger for other sites.
5. Spatial smoothing of over-prediction errors is possible through pairing of two spatially-distributed PV plants.

Additionally, in order to reduce the obstacles to safe and economic operation of solar power, research should shift focus towards power output forecasting. While improving irradiance prediction accuracy is important for many applications, our results suggest that for the day-ahead power forecasts, statistical methods can compensate for systematic errors in the irradiance predictions and are suitable for modeling the complexity of power plant operations.

2.6 Acknowledgements

This chapter is taken, in part, from: D. P. Larson, L. Nonnenmacher and C. F. M. Coimbra, “Day-Ahead Forecasting of Solar Power Output from Photovoltaic Plants in the American Southwest”, *Renewable Energy* (2016), 91, pp. 11–20. The dissertation author was the primary investigator and author of these manuscripts.

Chapter 3

Intra-Day Solar Power Forecasting

3.1 Prior Work

Most prior intra-day forecasting studies have focused on satellite-to-irradiance modeling together with techniques that use cloud motion vector (CMV) fields. Perez et al. (2010) predicted hourly GHI for 1–6 hours ahead at SURFRAD stations throughout the continental United States [PKS⁺10]. Marquez et al. (2013) developed a hybrid forecast model based on CMV fields and Artificial Neural Networks (ANNs) to predict GHI in Merced, CA, USA and Davis, CA, USA over horizons of 30–120 minutes ahead [MPC13]. Nonnenmacher and Coimbra (2014) forecasted GHI 1–3 hours ahead in San Diego, CA, USA using streamlines determined from CMV fields [NC14]. Zagouras et al. (2015) investigated the use of lagged exogenous variables and spatiotemporal correlations to improve 1–3 hour ahead GHI forecasts at seven CIMIS sites in California [ZPC15]. Most recently, Mazzora Aguiar et al. (2015) used ANNs to forecast GHI 1–6 hours ahead for sites in the Canary Islands, off the northwest coast of Africa [MAPD⁺15].

For a review of PO forecasting of solar power plants, readers are referred to Antonanzas et al. (2016) [AOU⁺16]. Here we highlight the PO forecasting literature that focused on sites in the American Southwest, the area of interest in this manuscript. Chu et al. (2015) evaluated intra-hour forecasts of PO of a 48 MW PV plant near Las Vegas, NV, USA [CUG⁺15]. Lipperheide et al.

(2015) forecasted the PO of the same 48 MW PV plant studied in Chu et al. (2015), but focused on horizons of less than 180 seconds ahead [LBK15]. Lonij et al. (2013) used a network of residential PV systems in Tucson, AZ, USA to forecast PO with horizons of 15–90 minutes [LBC⁺13]. Pedro and Coimbra (2012) evaluated 1–2 hour ahead PO forecasts, generated without exogenous inputs, for a 1 MWp PV farm in Merced, CA, USA [PC12]. Larson et al. (2016) forecast day-ahead PO of two 1 MWp PV plants in San Diego County, CA, USA [LNC16].

3.1.1 Contributions

We aim to present a forecast methodology to accurately and directly predict PO of PV power plants 1–6 hours ahead. Rather than forecast the solar resource and then develop a secondary resource-to-power (RTP) model, we focus on directly predicting the PO from the input variables. Removing the need for a RTP model minimizes data dependencies and increases the likelihood that our methodology will be adopted by utilities and power plant operators. Specifically, our methodology only requires historical PO and satellite imagery for training, whereas other methods may require meteorological data streams (humidity, temperature, *etc.*) as well as diagnostic information related to power plant operations (*e.g.* PV panel temperature) for the RTP model, in addition to the inputs for the solar resource forecast. Such data streams may not be available due to 1) costs of installation and maintenance or 2) information security requirements at the power plant.

Additionally, after training, our methodology does not depend on ground telemetry or PO data to generate forecasts. This means our forecast models are robust to issues at the target site, such as sensor malfunctions or measurement latency. There may be issues related to the satellite images, but a single data dependency (the satellite imagery) is preferable to a scenario where any one of multiple data streams could hinder or prevent generating a forecast. This is particularly important for power plant operators, who must make scheduling decisions with hard deadlines.

Our methodology is applied to two power plants in the American Southwest, a region undergoing rapid and large-scale growth in renewable energy generation. While forecast studies in other regions are still of importance and value, our results provide a valuable reference for parties




Figure 3.1: A diagram of the overall forecast methodology. First, a 480×680 pixel satellite image at time t is cropped down to a $w \times w$ region of interest, centered around the target site. The $w \times w$ cropped image is then transformed into a n -length feature vector $x(t)$ and fed into a forecast model, *e.g.*, SVR, to produce a prediction of the target output $\hat{y}(t + \Delta t)$.

involved in solar energy generation in the American Southwest, *e.g.*, power plant operators and utilities. The two plants are not the largest centralized power plants in the region, *e.g.*, the California Valley Solar Ranch (CVSR) near San Luis Obispo is 250 MW.. However, the size of the two sites (1 MWp each) is comparable to many existing and planned PV-based systems in the region, which lends confidence to the applicability of the results.

3.2 Forecast Methodology

We seek to learn a mapping $f(\cdot)$ between a set of real-valued satellite-derived features $x^{(i)} \in \mathbf{R}^n$ at the current time t and ground conditions $y^{(i)} \in \mathbf{R}$ at some future time $t + \Delta t$. More specifically, we are interested in finding $f : \mathbf{R}^n \mapsto \mathbf{R}$ for intra-day forecasting of PO of PV solar farms (see Figure 3.1). To accomplish these goals, we consider a linear model based on Ordinary Least-Squares (OLS) and a non-linear model based on Support Vector Regression (SVR).

Our model selections are influenced primarily by two goals. First, both researchers and non-researchers, *e.g.*, power plant operators, should be able to implement the models. OLS and SVR are standard regression methods, and are supported by most modern scientific programming environments, *e.g.*, MATLAB, Python and R. This increases the likelihood that our forecast methodology will be adopted and that our results can be reproduced. Second, there should be prior literature supporting the use of the models for solar forecasting tasks. Both OLS and SVR have been used successfully in published solar forecasting literature (see for example [KNPC16, FOO⁺14a, LNC16, WKL⁺16]).

To avoid ambiguity, we explicitly define several forecasting concepts as used in this study.

First, intra-day forecast are predictions with forecast horizons of 1–6 hours ahead. Second, the intra-day forecasts are for backwards-averaged hourly values. For example, suppose we wish to generate a set of intra-day forecasts at 9:00 for a site with per-minute data. The 1 hour ahead forecast generated at 9:00 would be a prediction of the hourly averaged value that covers the time span of 10:01–11:00. Then the 2 hour ahead forecast, also generated at 9:00, would be a prediction of the average over 11:01–12:00. And so on for the 3, 4, 5 and 6 hour ahead forecasts. Third, we define night as the period when the solar zenith angle (θ_z) is greater than 85 degrees. With these definitions, we move onto the forecast models.

3.2.1 Ordinary Least-Squares (OLS)

OLS is a linear regression method that solves the objective

$$\text{minimize } \sum_{i=1}^m \left(f(x^{(i)}; \boldsymbol{\theta}, b) - y^{(i)} \right)^2, \quad (3.1)$$

where $\{\boldsymbol{\theta} \in \mathbf{R}^n, b \in \mathbf{R}\}$ are the model parameters, and $f(x^{(i)}; \boldsymbol{\theta}, b) = \boldsymbol{\theta}^T x^{(i)} + b$ is the forecast function.

To discourage overfitting, we add a ℓ_2 regularization term to the OLS objective function

$$\text{minimize } \sum_{i=1}^m \left(f(x^{(i)}; \boldsymbol{\theta}, b) - y^{(i)} \right)^2 + \alpha \|\boldsymbol{\theta}\|_2, \quad (3.2)$$

where $\alpha \in \mathbf{R}$ is the ℓ_2 regularization constant¹.

3.2.2 Support Vector Regression (SVR)

SVR is a form of Support Vector Machines (SVM) used for non-linear regression tasks. In this study, we use v-SVR [Vap98, CL02], as implemented in LIBSVM [CL11]. The objective

¹OLS as described in Equation 3.2 is also known as Ridge Regression.

function is

$$\begin{aligned}
& \text{minimize} && \frac{1}{2} \|\boldsymbol{\theta}\|_2^2 + C \left(\nu \varepsilon + \frac{1}{m} \sum_{i=1}^m (\xi_i + \xi_i^*) \right) \\
& \text{subject to} && f(x^{(i)}; \boldsymbol{\theta}, b) - y^{(i)} \leq \varepsilon + \xi_i, \\
& && y^{(i)} - f(x^{(i)}; \boldsymbol{\theta}, b) \leq \varepsilon + \xi_i^*, \\
& && \xi_i, \xi_i^* \geq 0, i = 1, \dots, m, \\
& && \varepsilon \geq 0,
\end{aligned} \tag{3.3}$$

where $\{\boldsymbol{\theta} \in \mathbf{R}^n, b \in \mathbf{R}, \xi \in \mathbf{R}^m, \xi^* \in \mathbf{R}^m, \varepsilon \in \mathbf{R}\}$ are the model parameters, $f(x^{(i)}; \boldsymbol{\theta}, b) = \boldsymbol{\theta}^T \phi(x^{(i)}) + b$ is the forecast function, and $\phi : \mathbf{R}^n \mapsto \mathbf{R}^n$ maps the input $x^{(i)}$ to a higher dimensional feature space. The two remaining variables, $\nu \in (0, 1]$ and $C \in \mathbf{R}$, are hyperparameters. To improve computational performance, we use the radial basis function (RBF) kernel $K(\cdot)$

$$K(x, x') = \exp(\gamma \|x - x'\|_2^2), \tag{3.4}$$

where $\gamma \in \mathbf{R}$ is a hyperparameter.

3.2.3 Persistence

Following prior literature, we include a persistence forecast as a benchmark for evaluating our forecast methodology. Formally, we define the persistence forecast as

$$\hat{y}(t + \Delta t) = y(t), \tag{3.5}$$

where $y(t)$ is the target value at some time t , Δt is the forecast horizon, *e.g.*, 1 hour, and $\hat{y}(t + \Delta t)$ is the forecasted value.

3.2.4 Smart Persistence

We also include a smart persistence forecast as a secondary benchmark. The smart persistence model improves upon the persistence forecast by incorporating a clear sky model

$$\hat{y}(t + \Delta t) = \frac{y_{\text{clear}}(t + \Delta t)}{y_{\text{clear}}(t)} y(t), \quad (3.6)$$

where y_{clear} is the clear sky value. When the target (y) is GHI, the fraction in Equation 3.6 is commonly referred to as the clear sky index (k_t). We use the airmass independent model described in [IP02, Ine02, Gue12] to estimate clear sky GHI. For clear sky PO, we used approximately 160 manually identified clear days, covering all four seasons, to fit a linear mapping between clear sky GHI and PO for each PV plant. As discussed in Larson et al. (2016), the assumption of a linear relationship between GHI and PO is valid for the two sites [LNC16].

3.3 Data

We analyze our forecast methodology using two PV plants in the American Southwest: Canyon Crest Academy (CCA) in San Diego, CA and La Costa Canyon (LCC) in Carlsbad, CA. Both plants have non-tracking PV panels at a fixed 5° incline, with nameplate capacities of 1 MWp. Additionally, both CCA (32.959° , -117.190°) and LCC (33.074° , -117.230°) are less than 10 km from the Pacific Ocean (see Figure 3.2). However, the PV panels at LCC are aligned 30° southwest due to the site's characteristics, leading to a difference in peak output of the two sites under identical conditions. For further details on the two sites, readers are referred to [LNC16].




Figure 3.2: A sample visible channel image from the geosynchronous GOES-West satellite, with markers showing the two PV plants used in this study: CCA (blue solid circle) and LCC (green solid square). The satellite images are 8-bit grayscale, with orange overlay lines showing country, state and county borders.

3.3.1 Satellite Images

The Geostationary Operational Environmental Satellite (GOES) system is comprised of geosynchronous satellites operated by the National Oceanic and Atmospheric Agency (NOAA) through the National Environmental Satellite, Data and Information Service (NESDIS). For the purposes of this study, we will be only using visible wavelength images from GOES-15, which is currently designated as GOES-West. The visible image channel is centered at $0.63 \mu\text{m}$, with a spectral range of $0.53\text{--}0.75 \mu\text{m}$, a spatial resolution of 1 km, and a temporal resolution of one image per 30 minutes.

A $w \times w$ square region, centered on the target site, is extracted from each image and then flattened into a vector $\tilde{x}^{(i)} \in \mathbf{R}^n$, where $n = w \times w$. Each image is then normalized

$$x^{(i)} = \frac{\tilde{x}^{(i)} - \mathbf{avg}(\tilde{x}^{(i)})}{\sqrt{\mathbf{var}(\tilde{x}^{(i)}) + 10}}, \quad (3.7)$$




Figure 3.3: The merged PO and satellite image data sets for the two power plants: (a) CCA and (b) LCC. Two years of data (2012–2013) are designated for training the forecast models (shown in gray), with two additional years (2014–2015) used for testing (shown in black).

where $\tilde{x}^{(i)} \in [0, 255]^n$ is the unprocessed 8-bit grayscale image vector, $\mathbf{avg}(\cdot)$ is the arithmetic mean value, and $\mathbf{var}(\cdot)$ is the variance. In Equation 3.7, the numerator normalizes the brightness and the denominator normalizes the contrast. After normalization, the image vectors are stacked to create a matrix $X \in \mathbf{R}^{m \times n}$, where m is the number of images.

Four years (2012–2015) of hourly PO measurements and satellite-derive features are merged to create data sets for training and testing our intra-day PO forecast methodology (see Figure 3.3). Although the images are available every 30 minutes, we only use images from the top of each hour to generate the forecasts. Additionally, night values are removed from the data set, *i.e.*, only data points with $\theta_z < 85^\circ$ are included. The histograms in Figure 3.4 show that the training and testing sets have similar distributions of PO data.

The data sets used in our analysis exceed the overall time span of any known intra-day PO forecast studies. Pedro and Coimbra (2012) is the most analogous study to our manuscript due to its focus on PO from a 1 MWp PV plant [PC12]. However, their forecast testing set consisted of eight months, compared to the 24 months in our testing sets. And while the Pedro and Coimbra (2012) paper studied one 1 MWp PV plant, our work evaluates PO forecasts for two separate 1 MWp PV plants. However, this comparison is not meant to lessen the scientific value of Pedro and Coimbra (2012). Rather, the comparison is meant only to help understand the value of the data sets in our manuscript.




Figure 3.4: Histograms of PO [kW] in the training (left column) and testing (right column) data sets for CCA (top row) and LCC (bottom row), with night values removed. The lower peak PO of LCC compared to CCA is due to the site characteristics of LCC.

3.4 Results and Discussion

We evaluate our methodology by predicting PO at CCA and LCC. For each site, the forecasts are generated at the top of the hour using the most recent satellite image. In addition to the satellite-derived features, the OLS and SVR models are supplied clear sky values of PO at the forecast target time $t + \Delta t$. The clear sky values are included to provide the models with both diurnal and seasonal information. Ground telemetry is not included as input to the models in order to minimize data dependencies and improve the robustness of the models to ground-level operating conditions, *e.g.*, delay or loss of data streams.

The forecasting models are trained on two years of data (2012–2013) and then tested on a separate set of two years (2014–2015) to ensure unbiased performance evaluations. Cross-validation is used to select the hyperparameters of the OLS and SVR models for each horizon and site. Figure 3.3 illustrates the training and testing sets used for the four scenarios. The data sets shown are complete, *i.e.*, they show the union of the output (PO) and input variables (satellite images). Section 3.4.1 evaluates the effect of input image size choice on the forecast performance. In Section 3.4.2, we analyze the impact of the model choice, *i.e.*, OLS vs. SVR. And in Section 3.4.3, we compare the performance of our forecast methodology to prior intra-day forecasting literature.

3.4.1 Effect of Image Size

The input image size may impact forecast performance. To investigate the effects, we compare the RMSE of OLS and SVR intra-day forecast models trained on square image regions with $w \in \{4, 8, 16, 32, 64, 128\}$. As shown in Figure 3.5, increasing the input image size decreases the RMSE for both the OLS and SVR models in general for CCA. However, there are diminishing returns for both models as the number of features ($n = w \times w$) grows beyond $w = 64$. Doubling the image width from $w = 64$ to $w = 128$ increases the mean RMSE of the OLS models by 5.6% as well as the minimum and maximum RMSE values by 2% and 5%, respectively. For the SVR models, increasing from $w = 64$ to $w = 128$ decreases the maximum RMSE by 8% while causing a

Table 3.1: Forecast performance on the testing set (2014–2015) for CCA and LCC for horizons of 1–6 hours ahead. The column names (1h, ..., 6h) indicate the forecast horizon in hours. Negative MBE values correspond to models biased towards over-predicting PO and vice-versa. Skill is reported relative to a smart persistence model. Bolded values indicate the forecast with the best value for a specific error metric, site and forecast horizon.

	RMSE [kW]						MBE [kW]						s [%]					
	1h	2h	3h	4h	5h	6h	1h	2h	3h	4h	5h	6h	1h	2h	3h	4h	5h	6h
CCA																		
Pers.	283.8	390.1	473.0	525.9	557.1	551.6	40.0	81.4	125.4	175.6	204.8	228.4	—	—	—	—	—	—
Smart Pers.	137.1	201.5	268.5	329.2	376.9	401.9	51.6	95.0	140.3	185.0	219.8	243.6	—	—	—	—	—	—
OLS	114.4	126.1	134.6	136.9	135.0	128.4	3.2	4.9	2.2	-1.2	-2.2	-3.3	16.5	37.4	49.9	58.4	64.2	68.1
SVR	104.8	121.8	132.0	136.4	134.8	128.4	-8.5	-2.5	-13.3	-19.8	-21.9	-21.5	23.5	39.5	50.8	58.6	64.2	68.1
LCC																		
Pers.	232.1	316.6	383.5	428.8	457.3	456.4	33.5	69.8	109.5	154.7	183.7	205.3	—	—	—	—	—	—
Smart Pers.	117.9	173.7	233.4	289.1	333.7	359.6	47.8	87.2	129.4	171.6	205.4	227.2	—	—	—	—	—	—
OLS	96.6	106.5	114.5	118.1	117.1	113.2	-0.2	1.4	-1.6	-1.7	-1.2	-3.3	18.1	38.7	51.0	59.1	64.9	68.5
SVR	89.9	101.9	112.2	117.5	116.4	112.5	-10.3	-9.6	-13.7	-15.1	-15.4	-16.3	23.8	41.3	51.9	59.4	65.1	68.7




Figure 3.5: Effect of image size on forecast accuracy of the CCA training set (2012–2013), measured by RMSE [kW]. The mean RMSE for the OLS (solid) and SVR (dashed) models across the six horizons (1–6 hours ahead) are shown bounded by the minimum and maximum RMSE values. A similar trend is observed for the LCC training set.

negligible increase in mean RMSE (0.4%) and a non-negligible increase in the minimum RMSE, from 35 kW to 49 kW. A similar trend is observed for the LCC site. Hereafter, *all* forecast results are reported for models trained on 64×64 image regions.

The observed trends, *i.e.*, that $w = 64$ appears to be the optimal image size, can be explained by two complementary mechanisms. First, both OLS and SVR are well suited for data sets where the number of samples equals or exceeds the number of features, *i.e.*, $m = n$ or $m > n$. Images with $w = 64$ produce $n \sim 4k$ features, which is the same order of magnitude as the number of samples in the training sets for each forecast horizon ($m \sim 3k\text{--}6k$). In comparison, $w = 128$ images produce $n \sim 16k$ features, an order of magnitude larger than the number of samples, *i.e.*, $m \ll n$. Second, larger image sizes cover larger spatial regions, and therefore can provide the models with information on cloud systems further away from the target site. In other words, while the first mechanism provides pressure to decrease n relative to m , the second mechanism drives an increase in n . The balance between these two factors led to the optimal image width of $w = 64$.

3.4.2 Effect of Model Choice

Table 3.1 provides bulk statistics on the performance of the forecast models on the two years of testing data (2014–2015). For both sites, the SVR forecasts achieve the lowest RMSE and highest skill values for all horizons. For 1–2 hours ahead, the difference in skill between the OLS and SVR models is non-negligible ($\Delta s > 1\%$). However, beyond 3 hours, the two models achieve similar skill values ($\Delta s < 1\%$). This result implies that while SVR is the best of the models evaluated, both OLS and SVR are expressive enough to predict the intra-day PO behavior from satellite imagery. Additionally, for applications where SVR is too complex to implement, users of our forecast methodology can instead employ OLS and achieve comparable results.

For readers familiar with the forecast skill metric, the values in Table 3.1 may, at first glance, appear unusually large. Skill values reported in the PO forecasting literature rarely exceed 40%, regardless of forecast horizon (see [AOU⁺16]). However, the definition of forecast skill means the reference persistence model can have a large impact. *I.e.*, a poor performing reference model can inflate the performance of a forecast method, as measured by the skill. In this study, both the persistence and smart persistence models performed poorly (RMSE > 250kW) over the entire intra-day horizon range, thereby resulting in abnormally large, but nonetheless valid forecast skill values.

Besides bulk statistics, it is also important to consider error distributions when evaluating forecast models. As shown in Larson et al. (2016), violinplots present a compact solution to visualizing multiple forecast error distributions [HN98, LNC16]. Figure 3.8 uses violinplots to compare the error distributions of the OLS and SVR models across the forecast horizons (1–6 hours ahead). Together with the bulk statistics, the violinplots of Figure 3.8 reinforce the point that the OLS and SVR models are both expressive enough to predict the intra-day PO from the satellite imagery.

Table 3.2 summarizes forecast performance across all horizons (1–6 hours ahead), but breaks down the results by season: Spring (March–May), Summer (June–August), Fall (September–November), and Winter (December–February). Forecast performance is reported as relative RMSE

(rRMSE), which is the RMSE normalized by the mean PO (\overline{PO}) of the specified season. For all seasons and both sites, SVR achieves the lowest rRMSE values of the tested models. While the highest rRMSE values occur during the Winter for both sites (0.294 for CCA and 0.295 for LCC), the lowest rRMSE values are in the Fall for CCA (0.219) versus Summer for LCC (0.241). Additionally, while both sites employ identical PV technology and are separate spatially by less than 15 km, there is a non-negligible difference in forecast performance between the sites for Spring, Summer and Fall. One factor in this difference is the abundance of solar microclimates throughout the region (see [ZIC14]).

Figures 3.6 and 3.7 visualize the distributions of forecast absolute errors as functions of solar zenith and azimuthal angles. As we do not consider nighttime forecasts, there are fewer data points for the longer horizons. For example, the CCA testing set contains ~ 6800 1-hour ahead forecasts vs. ~ 3900 6-hour ahead forecasts. To prevent biasing the results, the plots only include 1-hour ahead forecasts. By grouping the errors by quartile, we see that the the largest errors, *i.e.*, the fourth quartile of errors, occur primarily in the morning. This result can be explained by two factors: 1) the high albedo in the visible channel satellite imagery near sunrise and 2) the prevalence of a marine layer in the region. Recent studies have proposed the use of other imaging wavelengths, *e.g.*, infrared, to reduce early morning forecast errors (see [HKWL15]). However, overlapping image datasets were not available at the time of this study and go beyond the scope of this work.

3.4.3 Comparison to Prior Work

The performance of our forecast methodology is comparable to prior intra-day PO studies. The Genetic Algorithm (GA) enhanced ANN from Pedro and Coimbra (2012) achieved skill values of 32.2% and 35.1%, relative to a smart persistence model, for horizons of 1 and 2 hours, respectively [PC12]. In comparison, our methodology achieves skill values of 30.6% (CCA) and 29.1% (LCC) for 1 hour ahead, and 46.1% (CCA) and 45.1% (LCC) for 2 hours ahead. The divergence in skill values for 2 hours ahead is possibly due to the GA-ANN model's lack of exogenous inputs, which was a self-imposed restriction by the authors [PC12]. Unfortunately, a lack




Figure 3.6: Kernel density estimates (KDE) of the absolute error (AE) of the 1-hour ahead SVR forecasts on the CCA testing set. The AE values are grouped into columns by quartile: (1) first quartile, (2) second quartile, (3) third quartile, and (4) fourth quartile. Darker areas represent higher densities. The top row (“Reference”) show the AE densities as functions of the solar azimuthal (ϕ) and zenith (θ_z) angles when the forecast was generated, while the bottom row (“Valid”) shows the solar angles of the targeted time for the forecasts. Sunrise and sunset are at $\theta_z = 90^\circ$ while $\phi < 180^\circ$ corresponds to the morning and $\phi > 180^\circ$ to the evening. The majority of the largest errors, *i.e.*, the fourth quartile of errors, are produced by forecasts generated near sunrise.




Figure 3.7: Same as Figure 3.6, but for the LCC testing set.

Table 3.2: Performance of forecasts across all studied horizons (1–6 hours ahead) on the testing set (2014–2015). Results are broken down by site and season: Spring (3/1–5/31), Summer (6/1–8/31), Fall (9/1–11/30), and Winter (12/1–2/28). RMSE values are reported relative to the mean PO of the season (\overline{PO}). Bold values indicate the forecast model with the lowest RMSE for a specific season.

rRMSE [–]	Spring	Summer	Fall	Winter	Total
CCA					
\overline{PO} [kW]	572.9	585.0	502.9	425.5	527.2
Pers.	0.857	0.849	0.868	0.892	0.868
Smart Pers.	0.516	0.533	0.532	0.586	0.539
OLS	0.227	0.239	0.226	0.300	0.244
SVR	0.223	0.232	0.219	0.294	0.238
LCC					
\overline{PO} [kW]	456.1	477.9	407.4	346.1	425.9
Pers.	0.869	0.855	0.888	0.908	0.880
Smart Pers.	0.541	0.594	0.592	0.641	0.590
OLS	0.250	0.250	0.250	0.296	0.259
SVR	0.243	0.241	0.244	0.295	0.252

of overlapping satellite imagery and PO data sets for the PV plant in Coimbra and Pedro (2012) prevents direct comparison.

Wolf et al. (2016) forecasted PO for 15 minutes–5 hours ahead for 921 PV systems of unspecified nameplate capacities [WKL⁺16]. Although not reported directly, approximate forecast skills can be derived from Figures 8 and 9 of Wolf et al. (2016). The SVR-based models achieved skill scores in the range of 20–36% for horizons of 1–5 hours ahead. Our methodology, which also uses SVR, achieves skill scores of 24–64% for the same forecast horizons. However, it is important to note that Wolf et al. (2016) used power data with a 15 minute temporal resolution while we forecasted hourly averaged values.

Although our work focuses on forecasting PO, comparisons to intra-day solar resource forecasting literature can still be of value. GHI forecast results are particularly well suited for comparison as the PO of both CCA and LCC are correlated with GHI due to their use of non-tracking PV panels (see [LNC16]). And fortunately, there have been multiple intra-day GHI forecasting studies at sites within the same region as both power plants. Nonnenmacher and Coimbra (2014) achieved GHI forecast skill scores of \sim 8–16% for horizons of 1, 2 and 3 hours at the University of

California San Diego campus (32.88, -117.23) [NC14]. Zagouras et al. (2015) obtained skills of $\sim 13\text{--}23\%$ (1–3 hours ahead) for the CIMIS weather station at Torrey Pines (32.90, -117.25) [ZPC15]. Perez et al. (2010) produced forecasts with skills of $\sim 6\text{--}19\%$ (1–6 hours ahead) for the SURFRAD station at the Desert Rock Airport, NV, USA (36.62, -116.02) [PKS⁺10]. The higher skill scores ($> 23\%$) achieved by our methodology, compared to these three analogous studies, indicates that our results are competitive with other intra-day forecast methods.

Another important comparison is between the intra-day PO forecast results of this manuscript and the day-ahead PO results of Larson et al. (2016) [LNC16]. Both papers use PO data from the same PV plants, *i.e.*, CCA and LCC, and cover similar time ranges (2011–2014 for Larson et al. (2016) vs. 2012–2015 for this study). Larson et al. (2016) reported day-ahead PO forecasts with skills of $\sim 15\text{--}23\%$ for CCA and $\sim 13\text{--}20\%$ for LCC, relative to a persistence model. In comparison, the intra-day PO forecasts attain higher skill scores than the day-ahead scenarios ($\sim 24\text{--}68\%$ for CCA and $\sim 24\text{--}69\%$ for LCC). Intra-day forecasts achieving higher skill scores than day-ahead forecasts is in agreement results from the forecasting literature, *e.g.*, Perez et al. (2010) [PKS⁺10].

3.5 Summary

We presented a methodology for forecasting PO of PV power plants over horizons of 1–6 hours ahead. Our methodology enables direct prediction of PO of operational solar farms from satellite imagery. This has two major benefits: First, by directly forecasting PO, we remove the need for intermediate irradiance forecasts and resource-to-power modeling, which are typically site specific. Second, by removing the need for additional data dependencies, *e.g.*, meteorological telemetry, we ensure that our methodology is generalizable to multiple power plants which may have varying levels of instrumentation. Together, these two features make our forecasting methodology well-posed to be applied to operational PV power plants.

We evaluated our forecast methodology using two non-tracking, 1 MWp PV plants in the American Southwest. Four years (2012–2015) of combined PO and satellite imagery data were




Figure 3.8: Forecast errors on the testing set (2014–2015) for (a) CCA and (b) LCC, grouped by horizon (1–6 hours ahead). The distributions for OLS (white) and SVR (blue) are shown side-by-side to enable direct comparisons of the two models for the same horizon. The horizontal dashed and dotted lines within each distribution represent the quartiles, with the width approximated using kernel density estimate (KDE) techniques. Following the the definition of MBE (Equation 2.9), negative error values correspond to the model over-predicting the PO.

used to analyze our methodology’s performance. Over a testing set of two years (2014–2015), our methodology achieved RMSE values of ~ 90 – 136 kW over horizons of 1–6 hours ahead and forecast skill scores of ~ 24 – 69% , relative to a baseline smart persistence model. Although there is limited prior intra-day PO forecasting literature, our results represent state-of-the-art performance, and validate the suitability of our methodology for intra-day PO forecasting.

3.6 Acknowledgements

This chapter is taken, in part, from: D. P. Larson and C. F. M. Coimbra “Direct Power Output Forecasts from Remote Sensing Image Processing”, *Journal of Solar Energy Engineering* (2018), 104 (2), pp. 021011-1–021011-08. The dissertation author was the primary investigator and author of this manuscript.

Chapter 4

Estimation of Cloud Optical Properties

Clouds play a pivotal role in radiative transfer through the atmosphere, but they are typically over-simplified in latency-constrained applications, e.g., solar forecasting. One factor in the over-simplification is the lack of real-time, accurate estimates of cloud optical properties, e.g., cloud optical depth (COD). Although the output of solar-based power generation systems, e.g., photovoltaic (PV) and concentrated solar power (CSP) plants, is strongly dependent on clouds and their optical properties, most solar forecasting methods rely on either binary or coarse representations of clouds, e.g., clear, partly cloud or overcast [MGC12, MPC13, CPN⁺14, CLPC15, LNC16, LC18, LCPC16, CUL⁺11]. In addition, data sources for estimated cloud optical properties (for a given location and time) tend to use quantized measures that are not consistent with physical quantities. For example, numerical weather prediction (NWP) models, e.g., the North American Mesoscale Forecast System (NAM), provide forecasts of cloud cover [%] per spatial grid point. While NWP forecasts of cloud cover have been used successfully as inputs for day-ahead (> 24 hours) solar power forecasting applications [LNC16], there is no generalized mapping from cloud cover to, e.g., COD. Therefore, the primary goal of this study is to provide a solution for real-time, accurate estimation of clouds and their optical properties that (1) can be applied to a large number of locations while (2) requiring minimal data dependencies.

Estimating physical properties from remote sensing is not a new nor completely solved

problem (see, e.g., [LMS01] or [WH12], and references therein). However, the increasing availability of data and computational resources, including software, present new opportunities. Specifically, the newest generation of geostationary satellites are opening up the possibility of high temporal (≤ 10 -minute) and spatial (≤ 2 km) resolution estimates, available concurrently across entire continents. In this work, we take a data-driven approach that couples remote sensing from the new geostationary satellites with physical modeling. While purely data-driven methods have seen recent success in many fields, e.g., Computer Vision, there is still value in coupling data with physical modeling. Particularly in the case where it is difficult, if not impossible, to derive a complete physical model of a phenomenon, but there exists an abundance of data that can be leveraged to complement the limitations of an incomplete or “coarse” physical model.

In addition, while estimates of physical properties from geostationary satellites are available, their implementations are typically skewed towards the needs of modeling, rather than latency-constrained applications. In a latency-constrained application, e.g., real-time solar forecasting, a moderately accurate estimate with low latency is more valuable than a high accuracy, but high latency estimate. For example, the National Oceanic and Atmospheric Administration (NOAA) provides data products for COD across the United States from its geostationary satellites. However, while the newest NOAA geostationary satellites have a 5-minute temporal resolution for the Continental United States (CONUS), the corresponding COD data product is only available every 15-minutes. Besides the increased latency, the reduced temporal resolution means less information is available for non-stationary clouds, which can have the largest impact on solar variability and therefore are the most important to forecast accurately. In this case, there is more value in providing a 5-minute estimate of COD, even if the estimate is less accurate than the existing 15-minute data product.

To meet the goal of providing low-latency estimates cloud optical properties, we present the Spectral Cloud Optical Property Estimation (SCOPE) method. SCOPE is based on solving the following constrained optimization problem:

$$\begin{aligned}
& \text{minimize} && f\left(\text{OLR}, \widehat{\text{OLR}}(x)\right) \\
& \text{subject to} && l \leq x \leq u,
\end{aligned} \tag{4.1}$$

where OLR is the outgoing longwave radiation at the top of the atmosphere (TOA), $\widehat{\text{OLR}}$ is the predicted OLR from a radiative model parameterized by a set of cloud optical properties $x \in \mathcal{X}$, $f(\cdot)$ is a statistical measure of the disagreement between the OLR measured by remote sensing and predicted by radiative modeling, and (l, u) are the lower and upper bounds on x , respectively.

4.1 Data

4.1.1 Remote sensing

The remote sensing data for the study comes from the latest generation of the Geostationary Operational Environmental Satellite (GOES) system: the GOES-R Series. Operated by the National Oceanic and Atmospheric Association (NOAA), the GOES-R Series consists of four identical satellites, each equipped with a suite of nadir-pointing, solar-pointing, and in-situ instruments (see Figure 4.1). Due to the focus on clouds, data from the nadir-pointing Advanced Baseline Imager (ABI) is considered. Compared to the imagers on the previous generation of GOES, the ABI provides data at five times the temporal resolution (Full Disk every 15 minutes and CONUS every 5 minutes), four times the spatial resolution (0.5–2.0 km), and three times the spectral resolution (16 spectral bands). Table 4.1 provides a summary of the ABI longwave ($\geq 4 \mu\text{m}$) spectral bands relevant to this study and Figure 4.2 shows the spectral response of the channels. Note that GOES-16, the first in the GOES-R Series, has been operational as GOES-East since December 2017, with the second in the GOES-R Series (GOES-17) only becoming operational as GOES-West in February 2019. Therefore, all analysis presented will be for GOES-16, but similar results are to be expected for GOES-17.

The GOES-R Series has a range of baseline products based on the ABI. We focus on




Figure 4.1: Illustration of the GOES-R Series satellite design, which includes a suite of instruments: the Advanced Baseline Imager (ABI), Geostationary Lightning Mapper (GLM), Advanced Baseline Imager (ABI), Solar Ultraviolet Imager (SUVI), Extreme Ultraviolet and X-ray Irradiance Sensors (EXIS), Space Environment In-Situ Suite (SEISS), and Magnetometer (MAG). The main satellite body has a volume of approximately $6\text{ m} \times 4\text{ m} \times 3\text{ m}$, while the five-panel solar array is approximately $7\text{ m} \times 4\text{ m}$. Figure adapted from <https://www.goes-r.gov/>.

the ABI Level 1b Radiance products for the CONUS scan. Specifically, the channel-averaged spectral radiance [$\text{W m}^{-2} \text{sr}^{-1} \text{cm}$] from the ten longwave channels (7–16), which have a spatial and temporal resolution of 2.0 km and 5-minutes, respectively. We retrieved a full year of data (1/1/2018–12/31/2018) from publicly-available sources and then processed each data file to extract the radiance per channel for the pixel corresponding to the target ground sites. Each ABI data file includes three UTC timestamps, corresponding to the start, center and end of the scan. To maintain compatibility with real-time applications, the end timestamp, rounded up to the next nearest 5-minute interval, is used to index the data.¹ The result is one CSV file per target site, consisting of eleven

¹For example, a scan with an end timestamp of 14:27 would be indexed as 14:30.

Table 4.1: Summary of the ten longwave channels from the GOES-R Series ABI, in both wavelength [μm] and wavenumber [cm^{-1}] basis. The center wavelengths (wavenumbers) are weighted-averages based on the integral of the spectral response function (SRF) of each channel, while the Full Width at Half Max (FWHM) wavelengths (wavenumbers) are the bandwidths corresponding to where the SRF exceeds 0.5 (50%). For a visualization of the channels and their SRFs, see Figure 4.2.

Channel	wavelength [μm]		wavenumber [cm^{-1}]	
	Center	FWHM	Center	FWHM
7	3.9	3.80–3.99	2570	2505–2630
8	6.2	5.79–6.59	1621	1517–1727
9	6.9	6.72–7.14	1444	1400–1487
10	7.3	7.24–7.43	1363	1346–1381
11	8.4	8.23–8.66	1184	1154–1216
12	9.6	9.42–9.80	1041	1020–1061
13	10.3	10.18–10.48	968	954–983
14	11.2	10.82–11.60	894	862–925
15	12.3	11.83–12.75	815	785–846
16	13.3	12.99–13.56	753	738–770

columns: one for the timestamp index and the remaining columns for the radiance measured by the longwave channels.

4.1.2 Ground telemetry

The Surface Radiation Budget Network (SURFRAD) provides long-term measurements of the surface radiation budget across the United States. There are seven active SURFRAD stations: Bondville, Illinois (BON); Desert Rock, Nevada (DRA); Fort Peck, Montana (FPK); Goodwin Creek, Mississippi (GWN); Pennsylvania State University, Pennsylvania (PSU); Sioux Falls, South Dakota (SXF); and Table Mountain, Boulder, Colorado (TBL). Table 4.2 summarizes the stations, while Figure 4.3 shows the locations of the stations from the perspective of GOES-16 and the terrain surrounding each station. Each station provides a range of irradiance and meteorological measurements at 1-minute resolution, along with quality control (QC) flags per variable. For this study, the following ground measurements are considered: ambient air temperature (T [K]), relative humidity (ϕ [%]), global horizontal irradiance (GHI [W m^{-2}]), direct normal irradiance

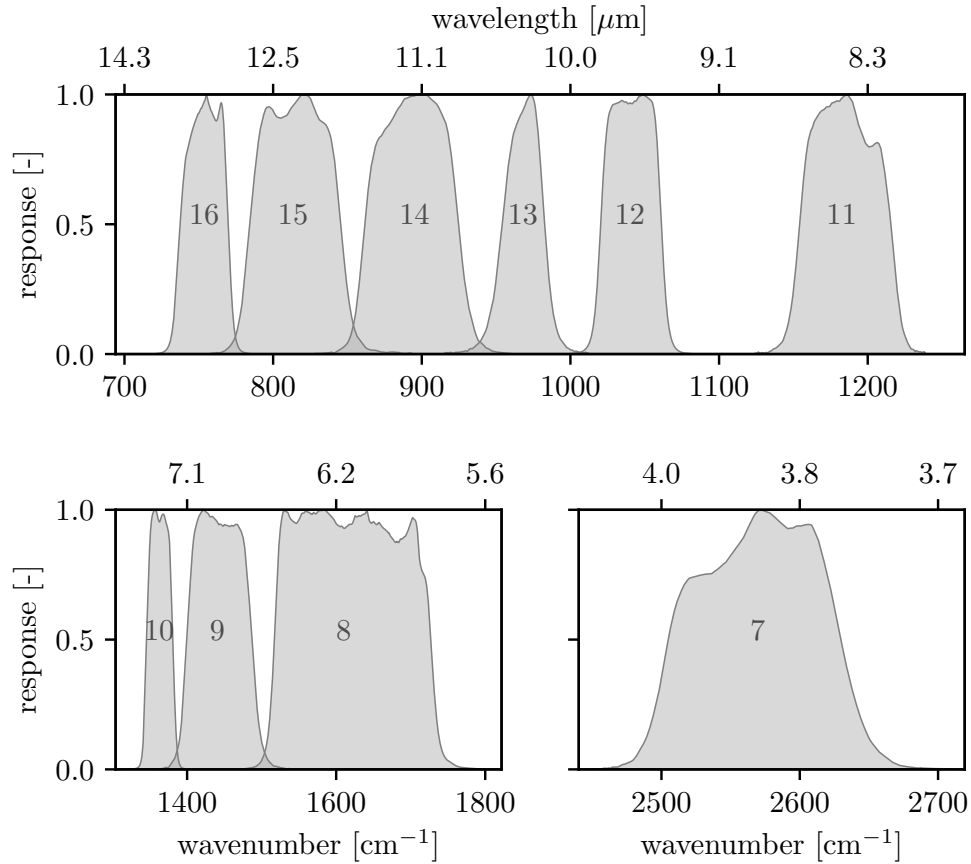


Figure 4.2: The spectral response function (SRF [-]) of the GOES-16 ABI longwave channels (7–16). For more details on the channels, see Table 4.1.

(DNI [W m^{-2}]), and downwelling longwave radiative flux (DLW [W m^{-2}]). All stations measure GHI with a Spectrolab SR-75 pyranometer ($0.28\text{--}3.0\ \mu\text{m}$), DNI with an Eppley Normal Incidence Pyrheliometer (NIP) ($0.25\text{--}3000\ \mu\text{m}$), DLW with an Eppley Precision Infrared Radiometer (PIR) ($3.0\text{--}50.0\ \mu\text{m}$), and temperature and humidity with a sensor located 10 m off the ground. We retrieved a full year of data (1/1/2018–12/31/2018) from each of the seven stations, removed questionable measurements (QC flags > 0), converted all variables to SI units, and then re-sampled the data to 5-minute backward-averaged values in order match the GOES-R CONUS scan rate. Note that GHI, DNI and DLW data are used as part of the analysis, rather than as inputs to the method.

Table 4.2: Summary of the seven SURFRAD stations: three-letter station code, latitude, longitude, and Köppen climate classification. The elevation (β) and azimuth (α) look angles are from the stations to GOES-16 (GOES-East), with $\alpha = 90^\circ$ and $\alpha = 180^\circ$ corresponding to due East and South, respectively. For a view of the stations from GOES-16, see Figure 4.3.

Station	Latitude [°]	Longitude [°]	Climate	β [°]	α [°]
BON	40.1	-88.4	Dfa	41.7	160.0
DRA	36.6	-116.0	Bwk	29.9	124.6
GWN	48.3	-105.1	Bsk	27.5	142.4
FPK	34.3	-89.9	Cfa	47.2	155.0
PSU	40.7	-77.9	Dfb	42.8	175.8
SXF	43.7	-96.6	Dfa	35.2	150.4
TBL	40.1	-105.2	Bsk	34.3	138.1

4.2 Methods

The purpose of the SCOPE method is to estimate cloud optical properties from longwave remote sensing data. The use of longwave spectral data has several benefits:

- Using longwave enables the method to be applied during day and night, thereby providing a unified approach to estimating cloud properties. While cloud information can improve daytime solar forecasts, nighttime cloud information can help improve solar forecasts for the often difficult to forecast early-morning hours.
- The longwave irradiance in the atmosphere can be approximated as diffuse, so there is no dependence on satellite view angles, solar angles and relative angles between the satellite, the sun and the ground sensors;
- Because of the diffuse nature of longwave irradiance, a two-flux radiative model [LJC17] can be used to estimate spectral longwave irradiance in the entire atmosphere at minimal computation cost;
- Most surfaces (grassland, desert, ocean) can be approximated as black surface in the longwave spectrum, while their shortwave albedo varies. Using longwave ensures reliable performance of our method across climatologically-diverse regions and during all seasons.

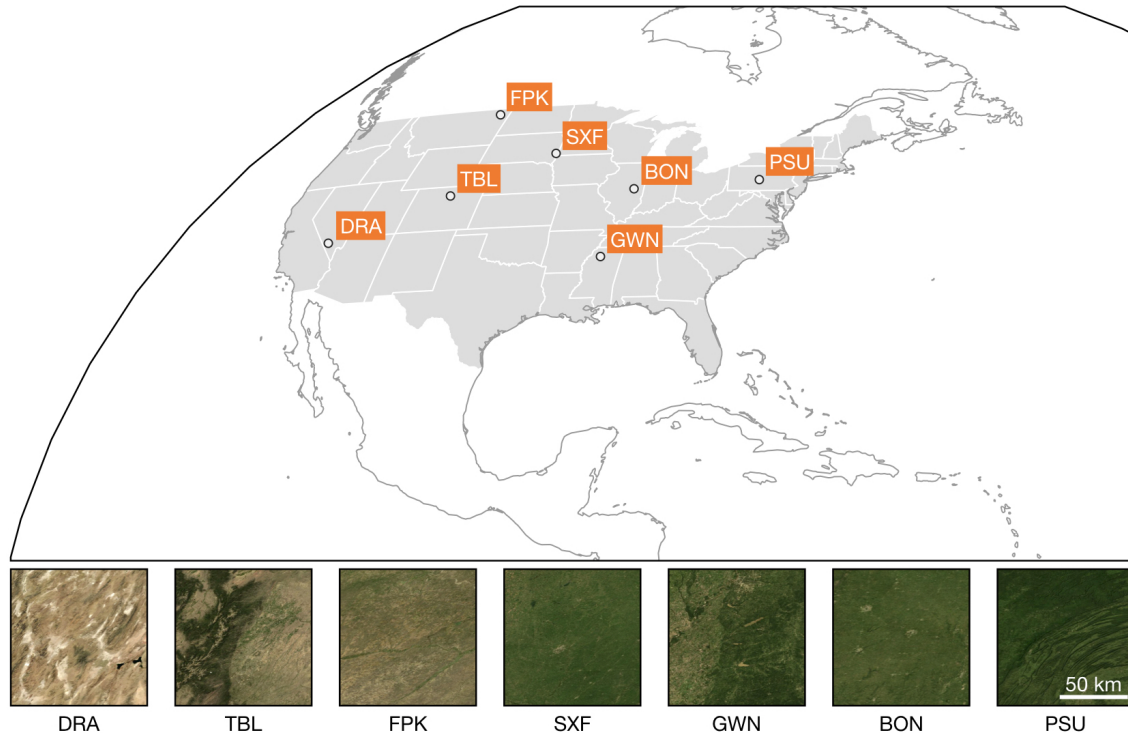


Figure 4.3: Top: The view of North America from the perspective of GOES-16, cropped to match the CONUS scan area. The US states are highlighted in solid grey, while the locations of the seven SURFRAD stations are denoted by circle markers. Bottom: satellite view of the $\sim 100 \text{ km} \times 100 \text{ km}$ region around each SURFRAD station, listed (left-to-right) from West to East. The satellite view highlights the climatological diversity of the stations, which cover arid desert (DRA), arid steppe (GWN, TBL), temperate (FPK), and continental (BON, PSU, SXF) climates. For more information on the SURFRAD stations, see Table 4.2.

- The outgoing shortwave irradiance measures the albedo of the atmosphere-Earth system, therefore it is generally difficult to identify and remove reflective snow/ice from shortwave remote sensing data, particularly in areas where the snow cover is non-uniform or varies over time.
- The outgoing longwave irradiance also provides information about cloud top temperature (i.e. cloud altitude) while outgoing shortwave irradiance is independent of cloud temperature.
- The longwave channels available from the GOES-16 ABI and similar instruments have clear connections to the physical properties of clouds and atmospheric processes. This opens up the potential for future extensions of the method to estimating other atmospheric parameters

concurrently.

4.2.1 SCOPE method

The SCOPE method estimates cloud optical properties by coupling remote sensing with radiative modeling of the OLR at TOA. For this study, we focus on estimating cloud optical depth (τ [-]) and cloud top height (z_T [m]). While the attenuation of DNI does not depend directly on the cloud temperature and therefore altitude, z_T is necessary for estimating cloud shadow position and has a direct effect on the DLW irradiance at the surface.

Given satellite measurements of the upwelling flux per discrete spectral channel i (F_i [W m^{-2}]), over a set of n spectral bands, and corresponding estimates of \hat{F}_i from a radiative model parameterized by x , the estimation task is formulated as a constrained optimization problem:

$$\begin{aligned} & \text{minimize} && \frac{1}{n} \sum_{i=1}^n f\left(F_i, \hat{F}_i(x)\right) && (4.2) \\ & \text{subject to} && l \leq x \leq u \end{aligned}$$

where $f(\cdot)$ is a statistical measure of the agreement between the measured and estimated F_i , and the box constraints ($l \leq x \leq u$) enforce bounds from the physics of the problem.

There are two complications with the above approach. First, the chosen source for the upwelling flux – the ABI Level 1B Radiance product – provides channel-averaged spectral intensity values ($I_{i,\nu}^\uparrow$ [$\text{W m}^{-2} \text{sr}^{-1} \text{cm}$]). However, using the SRF per channel (R_i [-]) from the ABI calibration, the F_i can be computed as:

$$F_i = \pi I_{i,\nu}^\uparrow \int R_i(\nu) d\nu. \quad (4.3)$$

Then, to maintain consistency, \hat{F}_i is computed as

$$\widehat{F}_i = \pi \int \widehat{I}_{i,\nu}^\uparrow(\nu) R_i(\nu) d\nu, \quad (4.4)$$

where \widehat{I}_ν is the modeled spectral intensity.

Second, as shown in Section 4.2.3, the range of F_i varies by more than an order of magnitude across the ten channels (0.005–0.450 W m⁻² cm). This variation in scale of F_i between channels may result in the channels with large F_i values biasing the τ and z_T estimates, thereby decreasing the accuracy of the method. To avoid this issue, the method uses normalized F_i values (F_i^* [-]), where the F_i are normalized by the blackbody emissive power per channel i ($E_{b,i}$ [W m⁻²]):

$$\begin{aligned} E_{b,i} &= \pi \int_{\nu_{i,1}}^{\nu_{i,2}} I_b(\nu, T) d\nu \\ &= \pi \int_{\nu_{i,1}}^{\nu_{i,2}} \frac{2hc^2\nu^3}{\exp\left(\frac{hc\nu}{k_B T}\right) - 1} d\nu, \end{aligned} \quad (4.5)$$

where I_b is the blackbody spectral intensity [W m⁻² sr⁻¹ cm], T is the measured ambient air temperature, $h = 6.626 \times 10^{-34}$ J s is Planck's constant, $c = 3 \times 10^8$ m s⁻¹ is the speed of light, and $k_B = 1.38 \times 10^{-23}$ J K⁻¹ is the Boltzmann constant.

4.2.2 Spectral radiative model

A key component of the proposed SCOPE method is the use of an efficient, spectrally-resolved radiative model. The method leverages the two-stream, line-by-line longwave radiative model from [LLC18]. For ease of discussion, the key characteristics and assumptions of the radiative model are highlighted: The atmosphere from the surface to 120 km is modeled as multiple plane-parallel layers, with the layer boundaries defined by a pressure coordinate system [LLC18]. The vertical temperature profile is adapted from the AFGL midlatitude summer profile, with linear corrections of troposphere temperature (0–14 km) using the air temperature measured at the surface

T , as shown in Figure 4.4. The vertical profiles of atmospheric gases are adapted from standard AFGL midlatitude profiles with corrections for current surface conditions, including water vapor, which is adjusted based on surface measurements of T and ϕ [LLC18, LC19]. The line-by-line absorption coefficients of atmospheric gases are retrieved from the HITRAN database [KGR⁺16] and MT_CKD continuum model [MPM⁺12]. The optical properties of aerosols are modeled by Mie theory with aerosol vertical profile adapted from CALIPSO measurements [LLC18]. The optical properties of plane-parallel clouds are modeled again by Mie theory with predefined droplet distribution (effective radius of $10\ \mu\text{m}$ and variance of 0.1) [LPC19]. Anisotropic scattering of aerosols and clouds are scaled as isotropic by δ -M approximation [LB82], which is found to be sufficiently accurate for the considered longwave spectral [LLC19]. With each layer being diffuse and isotropic after the scaling, a two-flux approach is used to recursively compute spectral upwelling and downwelling fluxes at all layer boundaries [LLC18]. From the upwelling flux at TOA, the F_i is then computed according to Equation (4.4).

4.2.3 Model tuning and sensitivity analysis

Before evaluating the accuracy of the proposed method, we first analyze the effects of key model parameters. The objective of the analysis is two-fold:

1. Tune the spectral model to meet the requirement of real-time estimation of cloud optical properties. This requires balancing model precision and accuracy with computation time, which are functions of number of atmospheric layers and spectral resolution. Given the 5-minute update schedule for the GOES-16 CONUS scans, we seek a model configuration with guaranteed run times under 1-minute using readily-available computing resources, e.g., a standard workstation-class machine with a 4-core CPU and 8 GB of RAM.
2. Understand and validate how the model responds to parameter changes related to the underlying physics. Specifically, the effects of changes to the surface conditions and cloud properties.

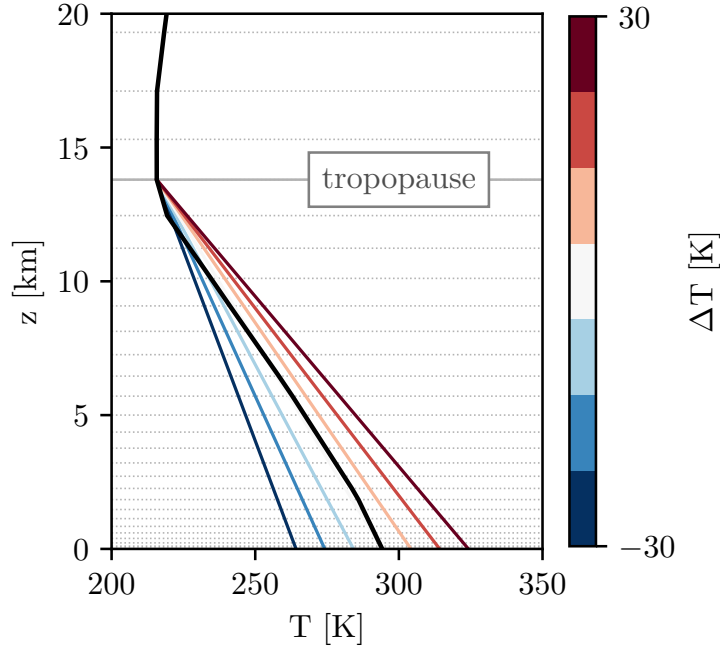


Figure 4.4: Illustration of the temperature profiles used in our method. The AFGL midlatitude summer profile (thick black line) is used as the base profile and then the temperature in the troposphere is linearly interpolated from the tropopause to the surface temperature. Here, the line colors correspond to ΔT , the difference between the original surface temperature provided by the AFGL profile (294.2 K) and the measured ambient temperature. The height (z) of the model layers (32-layer discretization) are shown as dotted lines. Over 2018, SURFRAD stations recorded temperatures between 233 K (-40°C) and 315 K (42°C), which correspond to ΔT values from -61 K to +21 K for the AFGL midlatitude summer profile.

The number of atmospheric layers (n [-]) and the spectral resolution ($\Delta\nu$ cm^{-1}) are selected by grid convergence test. The clear-sky broadband outgoing longwave irradiance (OLR) is computed as a function of n and $\Delta\nu$, with all other parameters constant. The AFGL midlatitude summer profile is selected to allow comparisons against prior work [LLC18]. Figure 4.5 and Table 4.3 present the results of the grid convergence test under clear-sky conditions. The results show that the model performance is more dependent on N than $\Delta\nu$, with the OLR changing by less than 0.1 W m^{-2} by going from $\Delta\nu = 0.10 \text{ cm}^{-1}$ to 0.01 cm^{-1} . In addition, going from 32 to 64 layers for $\Delta\nu = 0.10 \text{ cm}^{-1}$ changes the OLR by $< 1\%$. Therefore, a configuration of 32-layers and $\Delta\nu = 0.10 \text{ cm}^{-1}$ is selected, to balance between the model performance and computing times ($\ll 1$ -minute per run).

The spectral outgoing longwave irradiance with respect to the primary input parameters (T ,




Figure 4.5: Comparison of the clear-sky OLR as a function of the number of atmospheric layers (N) and spectral resolution ($\Delta\nu$) of the radiative model. All other model parameters are constant: $T = 294 \text{ K}$, $\phi = 50\%$, and $\tau = 0$. Top: the broadband OLR (integrated over $\nu \in (0, 3000] \text{ cm}^{-1}$) vs. N for $\Delta\nu \in \{0.01, 0.10, 1.00\} \text{ cm}^{-1}$. For each $\Delta\nu$, the model converges towards $\sim 278 \text{ W m}^{-2}$. Bottom: the absolute difference ($|\Delta|$) between the OLR of the three $\Delta\nu$ values. The $\Delta\nu = 0.01 \text{ cm}^{-1}$ data ends at $n = 64$ as the $n = 128$ case crashed due to insufficient memory.

ϕ , τ , and z_T) are analyzed using the radiative model to identify proper spectral channels for cloud estimation. Based on the underlying physics of radiative transfer in the atmosphere, the spectral intensity is dependent on the primary model parameters in the following ways as illustrated in Figure 4.6:

- Decreasing T causes the OLR to decrease as colder objects emit less longwave radiation.
- Increasing ϕ causes the OLR to decrease as water vapor attenuates longwave radiation emitted by the ground.
- Increasing τ increases cloud attenuation and reduces the objects of emission (replacing ground

Table 4.3: Convergence of the clear-sky broadband OLR [W m^{-2}] as the number of layers (n) increases, for $\Delta\nu \in \{1.00, 0.10, 0.01\} \text{ cm}^{-1}$. For each n_i , the change in OLR ($|\Delta|$) is reported relative to the OLR from the previous n_{i-1} . The symbol ** indicates cases that failed to run due to memory issues.

n	$\Delta\nu = 1.00 \text{ cm}^{-1}$		$\Delta\nu = 0.10 \text{ cm}^{-1}$		$\Delta\nu = 0.01 \text{ cm}^{-1}$	
	OLR [W m^{-2}]	$ \Delta $ [%]	OLR [W m^{-2}]	$ \Delta $ [%]	OLR [W m^{-2}]	$ \Delta $ [%]
4	295.611	–	296.217	–	296.272	–
8	285.962	3.26	286.495	3.28	286.617	3.26
16	280.156	2.03	280.874	1.96	280.903	1.99
32	278.173	0.71	279.260	0.57	279.253	0.59
64	277.077	0.39	278.239	0.37	278.254	0.36
128	276.739	0.12	277.583	0.24	**	**

with clouds), thereby decreasing the OLR.

- Increasing z_T causes the OLR to decrease as higher altitude clouds are colder and therefore emit less longwave radiation.

The relevant spectral ABI channels are readily identified:

- Channels 7 (shortwave window), 11 (cloud-top phase), and 13–15 (longwave windows) respond most strongly to changes in T and τ .
- Channels 8–10 (water vapor) respond the most to changes in ϕ .
- Channels 11 and 13–15 show the largest relative responses to z_T .
- Channels 8–10 respond to z_T , but are invariant to changes in τ .

Therefore, we select channels 8–10 to estimate z_T and then use channels 11 and 13–15 to estimate corresponding τ . The SCOPE method is detailed in Algorithm 1.

4.3 Results and Discussion

The accuracy of the proposed SCOPE method on four types of sky conditions is evaluated:

(1) daytime clear, (2) daytime cloudy, (3) daytime all-sky (clear or cloudy), and (4) nighttime all-sky.

Algorithm 1: Spectral Cloud Optical Property Estimation (SCOPE)

Input: $T, \phi, I_{7,v}^\uparrow, I_{8,v}^\uparrow, \dots, I_{16,v}^\uparrow$
Output: $\hat{\tau}, \hat{z}_T$

for $i \in \{7, \dots, 16\}$ **do**

$F_i = \pi I_{i,v}^\uparrow \int_{\nu_i} R_i(\nu) d\nu$	// flux per channel i
$E_{b,i} = \pi \int_{\nu_i} I_{b,v}(\nu, T) d\nu$	// blackbody flux
$F_i^* = F_i / E_{b,i}$	// normalized flux

end

$sky = \text{classify_sky}(T, \phi, I_{11,v}, I_{13,v}, I_{14,v}, I_{15,v})$ // estimate sky condition

if $sky == \text{cloudy}$ **then**

$\hat{z}_T = \arg \min_{z_T} \sum_{i \in \{8,9,10\}} f(F_i^*, \hat{F}_i^*(z_T; T, \phi))$	// estimate z_T
$\hat{\tau} = \arg \min_{\tau} \sum_{i \in \{7,11,13,14,15\}} f(F_i^*, \hat{F}_i^*(\tau, \hat{z}_T; T, \phi))$	// estimate τ

else

$\hat{\tau} = 0$	
$\hat{z}_T = 0$	

end

$\widehat{DLW} = \text{estimate_dlw}(T, \phi, \hat{\tau}, \hat{z}_T)$ // use estimated parameters to predict DLW

In all cases, the radiative model is configured with 32 layers, $\Delta\nu = 0.10 \text{ cm}^{-1}$, and use the same constant aerosol optical depth (AOD) as in Li et al. (2018) (AOD = 0.1243). [LLC18] For the daytime periods ($\theta_z \leq 85^\circ$), a clear-sky identification method, based on measured GHI, separates the data into daytime clear and daytime cloudy datasets. [IEC15] Over one full year (1/1/2018–12/31/2018), $\sim 15\%$ of the daytime data from all seven sites ($\sim 350,000$ samples total at 5-minute resolution) is identified as clear, with individual sites having $\sim 5\%$ (PSU) to $\sim 35\%$ (DRA) of their daytime data identified as clear.

Due to the lack of measurements of τ and z_T in the SURFRAD sites, the performance of the SCOPE method is analyzed using a two-part validation. First, the accuracy of the method at reproducing the OLR across the ten longwave channels is evaluated. If the τ and z_T estimates are correct, then the measured and modeled OLR should match for all ten channels. Second, the measured DLW [W m^{-2}] at the surface is compared with the DLW predicted by the radiative model (\widehat{DLW}) using the estimated τ and z_T . As no spectral response data is provided for the PIR instruments used to measure DLW at the stations, we compute \widehat{DLW} as

$$\widehat{\text{DLW}} = \pi \int_{\nu_1}^{\nu_2} \widehat{I}_{\downarrow}^{\nu}(\nu) d\nu, \quad (4.6)$$

where $\widehat{I}_{\downarrow}^{\nu}$ is the spectral downwelling radiance at the surface [$\text{W m}^{-2} \text{sr}^{-1} \text{cm}$], and ν_1 and ν_2 are the wavenumbers corresponding to the spectral band of the PIR (3–50 μm).

The key implementation details in our analysis are: all code is written in Python 3 using standard scientific Python libraries: numpy, scipy, and pandas. In addition, to minimize computational latency and therefore meet our goal of $\ll 60$ s per time instance, a look-up table of T , ϕ , τ and z_T combinations for F_i and DLW is pre-computed:

- $T \in \{234, 239, \dots, 319, 324\}$ [K],
- $\phi \in \{5, 10, \dots, 95, 100\}$ [%],
- $\tau \in \{0.10, 0.16, \dots, 63.1, 100.0\}$ [-],
- $z_T \in \{1.13, 1.84, 2.72, 3.76, 5.00, 6.46, 8.13\}$ [km],

where T is discretized in steps of 5 K from the AFGL midlatitude summer profile’s surface temperature (294 K), τ is discretized according to 10^x with $x \in \{-1.0, -0.8, \dots, 1.8, 2.0\}$, and z_T is discretized according to the pressure coordinate system of the radiative model. Note that the method does not explicitly select for the thickness of the cloud layers, nor does the method consider the case where multiple clouds, at different altitudes, overlap. The look-up table is used in place of directly calling the radiative model iteratively, thereby reducing the overall run time of the method. Specifically, at each time t , the modeled F_i are linearly interpolated to the $T(t)$ and $\phi(t)$ from ground measurements, and then τ and z_T are estimated from comparing modeled F_i with satellite measurements. On a standard workstation-class computer, the total wall-clock time for the estimation method using the look-up table is ~ 1 second per timestamp.

4.3.1 Daytime clear periods

First, the method is validated on clear-sky periods ($\tau = 0$) during daylight hours. Figure 4.7 provides a visualization of the F_i as a function of the water vapor partial pressure (P_w). The visualization only shows the results for one station (PSU), but all seven stations exhibit similar results. Namely, that the measured and modeled F_i^* agree across the range of observed P_w , with F_i decreasing as P_w increases, which corresponds to increased levels of water vapor content and therefore increased attenuation of longwave radiation. Figure 4.8 shows the RMSE and MBE of the F_i^* for the ten longwave channels and seven stations. Overall, the seven stations show similar levels of RMSE and MBE in each channel, with the largest error and variation in error among the stations in channel 7. Known as the shortwave window, channel 7 lies at the boundary of the shortwave-longwave spectral ranges and has a known reflected solar component during the day. As will be discussed in Section 4.3.4, the lower error in channel 7 during the nighttime all-sky case indicates that the higher channel 7 error is due to the reflected solar component. However, the higher channel 7 F_i^* error during the day is not sufficient to have a noticeable effect on the accuracy of the τ and z_T estimation, as measured by the DLW error.

To further validate the method during clear-sky periods, the DLW from the radiative model is compared against DLW measured by a PIR instrument at each station. Figure 4.9 compares the RMSE and MBE of the daytime clear-sky DLW at the stations. The DLW RMSE for five of the seven stations (BON, FPK, GWN, PSU, and SXF) is less than or approximately the same as the uncertainty of the PIR instrument ($\pm 5 \text{ W m}^{-2}$). The two stations with DLW RMSE $> 10 \text{ W m}^{-2}$ (DRA and TBL) are the furthest west of the seven stations and therefore at the most extreme look angles from GOES-16. In addition, the stations have the highest elevations of the seven (DRA: 1007 m, TBL: 1689 m). The extreme look angles and high elevations of DRA and TBL may be the cause of the large error, but further testing, e.g., of other sites with similar look angles and elevations, is required to provide a conclusive answer.

4.3.2 Daytime cloudy periods

The SCOPE method is then evaluated on daytime cloudy-sky periods. Here, both τ and z_T are estimated based on the F_i^* , and then validated using the DLW. Based on the spectral response of the F_i^* to τ and z_T , as discussed in Section 4.2.3, a two-step estimation method is adopted: First, z_T is estimated using channels 8–10, with a fixed COD value ($\tau = 0.1$). Then, using the estimated z_T value (\hat{z}_T), COD is estimated using channels 7, 11, 13–15. Compared to estimating τ and z_T concurrently, estimating z_T independent of COD allows us to project the COD estimation problem onto a smaller subspace, thereby reducing the number of potential solutions. Additionally, the two-step approach helps remove non-unique solutions that exist due to the underlying physics of the problem. For example, an optically thick, warm (low-altitude) cumulus cloud can result in the same F_i^* as an optically thin, cold (high-altitude) cirrus cloud.

Figures 4.8 and 4.9 visualize the OLR and DLW error, respectively, for the daytime cloudy-sky case. The RMSE and MBE of the F_i during cloudy-sky conditions is comparable to clear-sky conditions, with the highest error in channel 7. As expected, the RMSE and MBE of the cloudy-sky DLW is higher than the clear-sky as the τ and z_T values are estimated. In addition, the model under-predicts (MBE < 0) the cloudy-sky DLW for all sites. Under-predicting DLW means the estimation method is mostly likely under-predicting the cloud base height by assuming a pre-defined cloud thickness.

Note that in solar forecasting, high variability periods are the most difficult to predict. Therefore, the accuracy of the estimation method on clouds associated with high variability periods is more important than the overall accuracy. Typically, high variability periods are associated with clouds of low or medium value of COD, which do not fully attenuate the solar irradiance over a location. Therefore, the DLW error is compared as a function of the DNI clear-sky index ($k_t = \text{DNI}/\text{DNI}_{\text{clearsky}}$), which acts as a proxy for the optical thickness of the clouds due to the lack of cloud measurements at the SURFRAD stations. Figure 4.10 shows the distribution of DLW absolute errors, as a function of k_t , for PSU during daytime cloudy-sky periods. The largest DLW errors ($|\Delta| > 20 \text{ W m}^{-2}$) are primarily concentrated near $k_t < 0.1$, with the range $0.1 \leq k_t < 1.0$ have

comparably lower errors, thereby indicating that the estimation method performs best on optically thin clouds that do not attenuate the incoming solar irradiance entirely. To reinforce this conclusion, Figure 4.11 shows the DLW error metrics after removing samples where $k_t \leq 0.001$, where the DLW errors during daytime cloudy periods are reduced.

4.3.3 Daytime all-sky periods

The method is evaluated for daytime all-sky conditions, which includes both clear and cloudy conditions. First, to minimize errors from the method attempting to estimate clouds during clear-sky periods, a threshold-based heuristic is used to classify each timestamp as clear or cloudy. Based on empirical testing, a threshold of 0.05 on the mean error between the clear-sky modeled F_i^* ($\tau = 0$) and measured F_i^* for channels 11, 13–15 is sufficient for identify clear-sky periods for the purposes of this study. If a time instance is identified as clear, the method sets $\tau = 0$ and moves on to the next time step. Otherwise, the method estimates τ and z_T using the two-step approach detailed in Section 4.3.2. Figure 4.8 shows the F_i^* error per site, while Figures 4.9 and 4.11 shows the DLW error with and without clouds that strongly attenuate the incoming solar irradiance ($k_t < 0.001$). The F_i and DLW errors for the daytime all-sky case are similar to the daytime cloudy case, indicating that the proposed estimation method can be used for scenarios where ground-based identification of clear vs cloudy conditions are not available, e.g., sites without GHI data or during nighttime periods.

4.3.4 Nighttime all-sky periods

As a final test, the estimation method is evaluated on all-sky conditions at night ($\theta_z > 85^\circ$). The setup is the same as the daytime all-sky case (see Section 4.3.3), with one modification: we include a temperature inversion at 1 km in the AFGL midlatitude summer profile used by the radiative model. Based on empirical testing, the temperature inversion is necessary to achieve error metrics similar to the daytime cases. Note that while there are data sources for temperature profiles at night, e.g., atmospheric sounding, such data is not considered in order to minimize data dependencies and

therefore maximize the number of sites where the proposed method can be applied. Figures 4.8 and 4.9 show the F_i^* and DLW errors, respectively, for the nighttime all-sky periods. The nighttime all-sky error metrics are similar to the daytime all-sky case, with a few exceptions. First, the nighttime F_i^* error for channel 7 is lower than the daytime case, which is due to the channel 7 has a reflected solar component during the day. Second, the lack of DNI at night inhibits analysis of the effect of removing the thick clouds that most strongly attenuate radiation in the atmosphere. However, based on the change in DLW error from removing such clouds from the daytime cloudy and daytime all-sky cases (see Figures 4.9 and 4.11), we expect that the DLW errors for the nighttime all-sky case would decrease by a similar amount, e.g., RMSE from $\sim 30 \text{ W m}^{-2}$ to $\sim 20 \text{ W m}^{-2}$.

4.4 Summary

The SCOPE method for real-time, direct estimation of cloud optical properties from high-resolution longwave remote sensing data is presented. By considering only the longwave spectrum, the method provides a unified approach to estimating cloud properties during day or night. The SCOPE method builds upon a two-stream, spectrally-resolved radiative model, leverages publicly-available satellite imagery from GOES-R, and requires minimal ground telemetry as inputs (only ambient temperature and relative humidity), which can be provided by a standard weather station or numerical weather prediction forecasts. In addition, although SCOPE is evaluated on a GOES-R satellite (GOES-16), the method is compatible with other high-refresh satellites, e.g., the geostationary Himawari satellites that face the Eastern Hemisphere. SCOPE was evaluated on a full year of data (1/1/2018–12/31/2018) from all seven SURFRAD stations, for four cases: daytime clear-sky, daytime cloudy-sky, daytime all-sky (clear or cloudy), and nighttime all-sky. Results show that SCOPE can accurately estimate the τ and z_T of clouds during both daytime and nighttime, although the method achieves lower accuracy on optically thick clouds that fully attenuate DNI. Lastly, the Python-based implementation of the method has a total wall-clock run time of ~ 1 second per time instance t , which is sufficiently fast for latency constrained applications, such as solar irradiance and

power forecasting.

4.5 Acknowledgements

This chapter is taken, in part, from: D. P. Larson, M. Li and C. F. M. Coimbra “Direct Estimation of Cloud Optical Properties and Downwelling Longwave Radiation from GOES-R Spectral Imagery” (in preparation). The dissertation author is the primary investigator and author of this manuscript.




Figure 4.6: The upwelling spectral intensity at TOA (I_v^\uparrow [mW m⁻² sr⁻¹ cm]) as a function of (a) T at the surface, (b) ϕ at the surface, (c) τ , and (d) z_T , with all other parameters fixed as follows: (a) $\tau = 0$, $\phi = 50\%$; (b) $\tau = 0$, $T = 294$ K ($\Delta T = 0$ K); (c) $z_T = 3.2$ km (thickness: 0.5 km, cloud average temperature: 279 K), $T = 294$ K, $\phi = 50\%$; (d) $\tau = 0.1$, $T = 294$ K, $\phi = 50\%$. Note that the spectral intensity is shown with a resolution of 10 cm⁻¹ for improved readability, and that (a) shows the temperature difference from the AFGL midlatitude summer profile surface temperature ($\Delta T = T - T_0$, where $T_0 = 294$ K), rather than T , for ease of interpretation. For all plots, the location and bandwidth of the ten longwave ABI channels (7–16) are highlighted, with inset plots added to better show the the spectral range that contains channel 7 (2300–2700 cm⁻¹). The black dashed line in each subplot shows the blackbody spectral intensity ($I_{b,\nu}$) at $T = 294$ K.

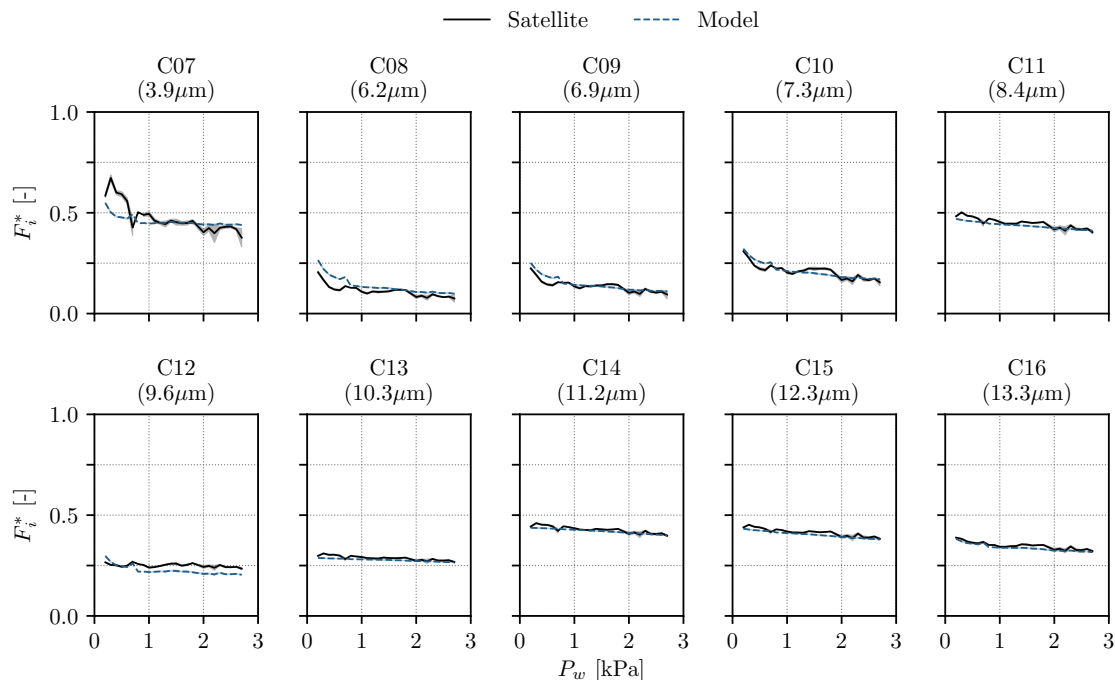


Figure 4.7: Comparison of the normalized flux per channel i (F_i^*) measured by the GOES-16 ABI (Satellite) and estimated by the longwave model (Model) for the ten longwave channels (7–16) versus the water vapor partial pressure (P_w [kPa]) at the surface. The data in the plots cover one full year (1/1/2018–12/31/2018) of daytime clear-sky periods for the PSU SURFRAD station. For each time t , the OLR is normalized by the blackbody radiation predicted by Planck’s distribution using the measured surface temperature ($T(t)$). Then the OLR are grouped into P_w bins of width 0.1 kPa, with the mean values show as the lines and the solid fill-in areas showing the 95% confidence interval (CI). Higher P_w values correlate with higher concentrations of water vapor in the air.




Figure 4.8: The RMSE and MBE between the measured and modeled normalized flux per channel (F_i^* [-]) for the ten longwave channels (7–16), seven SURFRAD stations, and the four evaluated cases: (1) daytime clear-sky, (2) daytime cloudy-sky, (3) daytime all-sky (clear or cloudy), and (4) nighttime all-sky (clear or cloudy). The division between daytime and nighttime is based on the solar zenith angle, with the clear versus cloudy periods segmented using ground measurements of GHI. For the daytime clear-sky case, COD is set to zero, whereas for the other three cases, τ and z_T are estimated at each time step from the F_i^* . The higher error in channel 7 (shortwave window) for all sites during the daytime can be attributed to the fact that channel 7 contains a reflected solar component during the day. For more information, see the GOES-R ABI documentation: <https://www.goes-r.gov/>.




Figure 4.9: The RMSE and MBE between the measured and modeled DLW [W m^{-2}] for the seven SURFRAD stations and four evaluated cases: (1) daytime clear-sky, (2) daytime cloudy-sky, (3) daytime all-sky (clear or cloudy), and (4) nighttime all-sky (clear or cloudy). For the daytime cloudy-sky, daytime all-sky and nighttime all-sky cases, the DLW is modeled using the τ and z_T estimated from the OLR. Negative MBE values ($\text{MBE} < 0$) correspond to the model under-predicting DLW, which can be attributed to the estimation method under-predicting τ , over-predicting z_T or a combination of the two.




Figure 4.10: The distribution of DLW absolute error ($|\Delta| := |\text{DLW} - \widehat{\text{DLW}}|$) for daytime cloudy-sky periods at PSU as a function of the DNI clear-sky index: $k_t = \text{DNI}/\text{DNI}_{\text{cs}}$, where DNI_{cs} is the clear-sky DNI. Darker colors correspond to $(k_t, |\Delta|)$ combinations with higher frequency of occurrence in the dataset ($\sim 45,000$ samples for 2018).




Figure 4.11: Same as Figure 4.9, but with the RMSE and MBE computed after removing samples with clouds that fully attenuate the incoming solar irradiance, using $k_t \leq 0.001$ as a threshold.

Chapter 5

Conclusions

The global transition to a low-carbon electric grid brings with it a range of technological, societal and policy challenges. One key technological challenge is the integration of intermittent renewable resources, e.g., solar and wind. The inherently variable and non-dispatchable nature of solar and wind-based power generation requires accurate forecasts for a range of horizons, from minutes to hours to days ahead. And as the penetration of renewable resources continue to grow, so does the value and importance of forecasting. In this work, we presented contributions in two key areas related to solar forecasting: (1) the development of data-driven methods to directly forecast the generation of solar-based power plants and (2) a novel method for providing real-time, accurate estimates of cloud optical properties, which can be used for modeling and forecasting applications.

The first half of the dissertation focused on data-driven approaches to forecasting the output of solar power plants. We presented methods that enable directly forecasting power output from data for both intra-day (>1 -hour) and day-ahead (>24 -hour) horizons. The methods remove the need for intermediate solar irradiance forecasts and resource-to-power modeling. In addition, the methods have minimal data dependencies and leverage readily-available public data sources, thereby helping to maximize the number of power plants where the methods can be applied. For both the intra-day and day-ahead horizons, the performance of the methods is validated using multiple years of generation data from two operational, 1 MW photovoltaic power plants in the American Southwest.

The results indicate the methods provide accurate solar power forecasts, and are robust to both seasonal and interannual variations.

In the remainder of the dissertation, we presented a novel method for providing real-time, accurate estimates of cloud optical properties directly from satellite imagery. While the output of solar-based power generation systems is strongly dependent on clouds, clouds and their optical properties are usually over-simplified due to a lack of low-latency, accurate estimates. We detailed an estimation method that builds upon a two-stream, spectrally-resolved radiative model coupled with high-resolution (5-minute, 2 km) spectral satellite imagery from GOES-R, the newest generation of the Geostationary Observational Environmental Satellite (GOES) system. The method relies solely on longwave radiation and therefore provides a unified approach to estimating cloud optical properties during both the day and night. We performed a two-step validation of the proposed method against a full year of data (1/1/2018–12/31/2018) using data from seven high-quality surface radiation observatories spread across climatologically diverse regions of the United States. The results show that the proposed method can provide real-time (~ 1 second per sample), accurate estimates of cloud optical depth (COD) and cloud top height during by daytime and nighttime all-sky (clear or cloudy) conditions.

Bibliography

- [AOU⁺16] J. Antonanzas, N. Osorio, R. Urraca, F. J. Martinez de Pison, and F. Antonanzas-Torres. Review of photovoltaic power forecasting. *Solar Energy*, 136:78–111, 2016.
- [APN15] M. P. Almeida, O. Perpignan, and L. Narvarte. PV power forecasting using a nonparametric PV model. *Solar Energy*, 115:354–368, 2015.
- [BMN09] P. Bacher, H. Madsen, and H. A. Nielsen. Online short-term solar power forecasting. *Solar Energy*, 83:1772–1783, 2009.
- [BRAH14] D. Bernecker, C. Riess, E. Angelopoulou, and J. Hornegger. Continuous short-term irradiance forecasts using sky images. *Solar Energy*, 110:303–315, 2014.
- [CA08] M. Chaabene and M. B. Ammar. Neuro-fuzzy dynamic model with Kalman filter to forecast irradiance and temperature for solar energy systems. *Renewable Energy*, 33:1435–1443, 2008.
- [CL02] C. C. Chang and C. J. Lin. Training ν -Support Vector Regression: theory and algorithms. *Neural Computation*, 14:1959–1977, 2002.
- [CL11] C. C. Chang and C. J. Lin. LIBSVM: A library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011.
- [CLPC15] Y. Chu, M. Li, H. T. C. Pedro, and C. F. M. Coimbra. Real-time forecasting of solar irradiance ramps with smart image processing. *Solar Energy*, 114:91–104, 2015.
- [CPC13] Y. Chu, H. T. C. Pedro, and C. F. M. Coimbra. Hybrid intra-hour DNI forecasts with sky image processing enhanced by stochastic learning. *Solar Energy*, 98:592–603, 2013.
- [CPN⁺14] Y. Chu, H. T. C. Pedro, L. Nonnenmacher, R. H. Inman, Z. Liao, and C. F. M. Coimbra. A smart image-based cloud detection system for intra-hour solar irradiance forecasts. *Journal of Atmospheric and Oceanic Technology*, 31(9):1995–2007, 2014.
- [CUG⁺15] Y. Chu, B. Urquhart, S. M. I. Gohari, H. T. C. Pedro, J. Kleissl, and C. F. M. Coimbra. Short-term reforecasting of power output from a 48 mwe solar pv plant. *Solar Energy*, 112:68–77, 2015.

- [CUL⁺11] C. W. Chow, B. Urquhart, M. Lave, A. Dominguez, J. Kleissl, J. Shields, and B. Washom. Intra-hour forecasting with a total sky imager at the uc san diego solar energy testbed. *Solar Energy*, 85(11):2881–2893, 2011.
- [CYL14] H. Y. Cheng, C. C. Yu, and S. J. Lin. Bi-model short-term solar irradiance prediction using support vector regressors. *Energy*, 70:121–127, 2014.
- [DFM15] M. Delfanti, D. Falabretti, and M. Merlo. Energy storage for PV power plant dispatching. *Renewable Energy*, 80:61–72, 2015.
- [DLM15] A. Dolara, S. Leva, and G. Manzolini. Comparison of different physical models for PV power output prediction. *Solar Energy*, 119:83–99, 2015.
- [FOO⁺14a] J. G. Fonseca, T. Oozeki, H. Ohtake, K. Shimose, T. Takashima, and K. Ogimoto. Regional forecasts and smoothing effect of photovoltaic power generation in Japan: an approach with principal component analysis. *Renewable Energy*, 68:403–413, 2014.
- [FOO⁺14b] J. G. Fonseca, T. Oozeki, H. Ohtake, T. Takashima, and K. Ogimoto. Regional forecasts of photovoltaic power generation according to different data availability scenarios: a study of four methods. *Progress in Photovoltaics: Research and Applications*, 2014.
- [Gue12] C. A. Gueymard. Clear-sky irradiance predictions for solar resource mapping and large-scale applications: improved validation methodology and detailed performance analysis of 18 broadband radiative models. *Solar Energy*, 86(8):2145–2169, 2012.
- [GUY⁺14] M. I. Gohari, B. Urquhart, H. Yang, B. Kurtz, D. Nguyen, C.W. Chow, M. Ghonima, and J. Kleissl. Comparison of solar power output forecasting performance of the total sky imager at the University of California, San Diego sky imager. *Energy Procedia*, 49:2340–2350, 2014. Proceedings of the SolarPACES 2013 International Conference.
- [GW11] C. A. Gueymard and S. M. Wilcox. Assessment of spatial and temporal variability in the US solar resource from radiometric measurements and predictions from models using ground-based or satellite data. *Solar Energy*, 85:1068–1084, 2011.
- [HKNF14] R. Hanna, J. Kleissl, A. Nottrott, and M. Ferry. Energy dispatch schedule optimization for demand charge reduction using a photovoltaic-battery storage system with solar forecasting. *Solar Energy*, 103:269–287, 2014.
- [HKWL15] A. Hammer, J. Kuhnert, K. Weinreich, and E. Lorenz. Short-term forecasting of surface solar irradiance based on Meteosat-SEVIRI data using a nighttime cloud index. *Remote Sensing*, 7:9070–9090, 2015.
- [HN98] J. L. Hintze and R. D. Nelson. Violin plots: a box plot-density trace synergism. *The American Statistician*, 52(2):181–184, 1998.

- [IEC15] R. H. Inman, J. G. Edson, and C. F. M. Coimbra. Impact of local turbidity estimation on clear sky models for direct normal irradiance. *Solar Energy*, 117:125–138, 2015.
- [Ine02] P. Ineichen. Comparison of eight clear sky broadband models against 16 independent data banks. *Solar Energy*, 80(4):468–478, 2002.
- [IP02] P. Ineichen and R. Perez. A new airmass independent formulation for the Linke turbidity coefficient. *Solar Energy*, 73(3):151–157, 2002.
- [IPC13] R. H. Inman, H. T. C. Pedro, and C. F. M. Coimbra. Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science*, 6:535–576, 2013.
- [KGR⁺16] R. V. Kochanov, I. E. Gordon, L. S. Rothman, P. Wcisło, C. Hill, and J. S. Wilzewski. HITRAN application programming interface (HAPI): A comprehensive approach to working with spectroscopic data. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 177:15–30, 2016.
- [KNPC16] A. Kaur, L. Nonnenmacher, H. T. .C. Pedro, and C. F. M. Coimbra. Benefits of solar forecasting for energy imbalance markets. *Renewable Energy*, 86:819–830, 2016.
- [KSHM13] B. Kraas, M. Schroedter-Homscheidt, and R. Madlener. Economic merits of a state-of-the-art concentrating solar power forecasting system for participation in the spanish electricity market. *Solar Energy*, 93:244–255, 2013.
- [LB82] H. Lee and R. O. Buckius. Scaling anisotropic scattering in radiation heat transfer for a planar medium. *Journal of Heat Transfer*, 104(1):68–75, 1982.
- [LBC⁺13] V. P. A. Lonij, A. E. Brooks, A. D. Cronin, M. Leuthold, and K. Koch. Intra-hour forecasts of solar power production using measurements from a network of irradiance sensors. *Solar Energy*, 97:58–66, 2013.
- [LBK15] M. Lipperheide, J. L. Bosch, and J. Kleissl. Embedded nowcasting method using cloud speed persistence for a photovoltaic power plant. *Solar Energy*, 112:232–238, 2015.
- [LC18] D. P. Larson and C. F. M. Coimbra. Direct power output forecasts from remote sensing image processing. *J. Sol. Energy Eng*, 104(2):021011–1–021011–08, 2018.
- [LC19] M. Li and C. F. M. Coimbra. On the effective spectral emissivity of clear skies and the radiative cooling potential of selectively designed materials. *International Journal of Heat and Mass Transfer*, 135:1053–1062, 2019.
- [LCPC16] M. Li, Y. Chu, H. T. C. Pedro, and C. F. M. Coimbra. Quantitative evaluation of the impact of cloud transmittance and cloud velocity on the accuracy of short-term dni forecasts. *Renewable Energy*, 86:1362–1371, 2016.

- [LFRAPV⁺12] V. Lara-Fanego, J.A. Ruiz-Arias, D. Pozo-Vázquez, F.J. Santos-Alamillos, and J. Tovar-Pescador. Evaluation of the WRF model solar irradiance forecasts in Andalusia (southern Spain). *Solar Energy*, 86(8):2200–2217, 2012.
- [LH74] A. Lacis and J. Hansen. A parameterization for the absorption of solar radiation in the Earth’s atmosphere. *Journal of Atmospheric Sciences*, 31(1):118–133, 1974.
- [LHHB09] E. Lorenz, J. Hurka, D. Heinemann, and H. G. Beyer. Irradiance forecasting for the power prediction of grid-connected photovoltaic systems. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2(2):2–10, 2009.
- [LJC17] M. Li, Y. Jiang, and C. F. M. Coimbra. On the determination of atmospheric longwave irradiance under all-sky conditions. *Solar Energy*, 144:40–48, 2017.
- [LLC18] M. Li, Z. Liao, and C. F. M. Coimbra. Spectral model for clear sky atmospheric longwave radiation. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 209:196–211, 2018.
- [LLC19] Z. Liao, M. Li, and C. F. M. Coimbra. Anisotropic corrections for the downwelling radiative heat transfer flux from various types of aerosols. *International Journal of Heat and Mass Transfer*, 136:1006–1016, 2019.
- [LMK14] J. Luoma, P. Mathiesen, and J. Kleissl. Forecast value considering energy pricing in california. *Applied Energy*, 125:230–237, 2014.
- [LMS01] J. Li, W. P. Menzel, and A. J. Schreiner. Variational retrieval of cloud parameters from goes sounder longwave cloudy radiance measurements. *Journal of Applied Meteorology*, 40(3):312–330, 2001.
- [LNC16] D. P. Larson, L. Nonnenmacher, and C. F. M. Coimbra. Day-ahead forecasting of solar power output from photovoltaic plants in the American Southwest. *Renewable Energy*, 91:11–20, 2016.
- [LPC19] M. Li, H. B. Peterson, and C. F. M. Coimbra. Radiative cooling resource maps for the contiguous united states. *Journal of Renewable and Sustainable Energy*, submitted, 2019.
- [LPKT14] E. W. Law, A. A. Prasad, M. Kay, and R. A. Taylor. Direct normal irradiance forecasting and its application to concentrated solar thermal output forecasting – a review. *Solar Energy*, 108:287–307, 2014.
- [LSH⁺11] Elke Lorenz, Thomas Scheidsteger, Johannes Hurka, Detlev Heinemann, and Christian Kurz. Regional PV power prediction for improved grid integration. *Progress in Photovoltaics: Research and Applications*, 19(7):757–771, 2011.
- [LSS14] Y Li, Y Su, and L Shu. An ARMAX model for forecasting the power output of a grid connected photovoltaic system. *Renewable Energy*, 66:78–89, 2014.

- [MAPD⁺15] L. Mozorra Aguiar, B. Pereira, M. David, F. Diaz, and P. Lauret. Use of satellite data to improve solar radiation forecasting with Bayesian Artificial Neural Networks. *Solar Energy*, 122:1309–1324, 2015.
- [MBCCM⁺14] D. Masa-Bote, M. Castillo-Cagigal, E. Matallanas, E. Caamano-Martin, A. Gutierrez, F. Monasterio-Huelin, and J. Jimenez-Leube. Improving photovoltaics grid integration through short time forecasting and self-consumption. *Applied Energy*, 125:103–113, 2014.
- [MC11] R. Marquez and C. F. M. Coimbra. Forecasting of global and direct solar irradiance using stochastic learning methods, ground experiments and the NWS Database. *Solar Energy*, 85:746–756, 2011.
- [MC13] R. Marquez and C. F. M. Coimbra. Intra-hour DNI forecasting methodology based on cloud tracking image analysis. *Solar Energy*, 91:327–336, 2013.
- [MCK13] P. Mathiesen, C. Collier, and J. Kleissl. A high-resolution, cloud-assimilating numerical weather prediction model for solar irradiance forecasting. *Solar Energy*, 92:47–61, 2013.
- [MGC12] R. Marquez, V. Gueorguiev, and C. F. M. Coimbra. Forecasting of global horizontal irradiance using sky cover indices. *J. Sol. Energy Eng*, 135:0110171–0110175, 2012.
- [MK11] P. Mathiesen and J. Kleissl. Evaluation of numerical weather prediction for intra-day solar forecasting in the continental United States. *Solar Energy*, 85(5):967–977, 2011.
- [MP10] A. Mellit and A. M. Pavan. A 24-h forecast of solar irradiance using artificial neural network: Application for performance prediction of a grid-connected pv plant at trieste, italy. *Solar Energy*, 84:807–821, 2010.
- [MPC13] R. Marquez, H. T. C. Pedro, and C. F. M. Coimbra. Hybrid solar forecasting method uses satellite imaging and ground telemetry as inputs to ANNs. *Solar Energy*, 92:176–188, 2013.
- [MPM⁺12] E. J. Mlawer, V. H. Payne, J.-L. Moncet, J. S. Delamere, M. J. Alvarado, and D. C. Tobin. Development and recent evaluation of the MT_CKD model of continuum absorption. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 370(1968):2520–2556, 2012.
- [MZP⁺10] L. Martin, L. F. Zarzalejo, J. Polo, A. Navarro, R. Marchante, and M. Cony. Prediction of global solar irradiance based on time series analysis: Application of solar thermal power plants energy production planning. *Solar Energy*, 84:1772–1781, 2010.
- [NC14] L. Nonnenmacher and C. F. M. Coimbra. Streamline-based method for intra-day solar forecasting through remote sensing. *Solar Energy*, 108:447–459, 2014.

- [NKC14] L. Nonnenmacher, A. Kaur, and C. F. M. Coimbra. Verification of the SUNY direct normal irradiance model with ground measurements. *Solar Energy*, 99:246–258, 2014.
- [NKC15] L. Nonnenmacher, A. Kaur, and C. F. M. Coimbra. Day-ahead resource forecasting for concentrated solar power integration. *Renewable Energy*, 86:866–876, 2015.
- [OSF⁺13] H. Ohtake, K. I. Shimose, J. G. D. S. Fonseca, T. Takashima, T. Oozeki, and Y. Yamada. Accuracy of the solar irradiance forecasts of the Japan Meteorological Agency mesoscale model for the Kanto region, Japan. *Solar Energy*, 98, Part B:138–152, 2013.
- [PC12] H. T. C. Pedro and C. F. M. Coimbra. Assessment of forecasting techniques for solar power output with no exogenous inputs. *Solar Energy*, 86:2017–2028, 2012.
- [PC15] H. T. C. Pedro and C. F. M. Coimbra. Short-term irradiance forecastability for various solar micro-climates. *Solar Energy*, 122:587–602, 2015.
- [PGK13] S. Pelland, G. Galanis, and G. Kallos. Solar and photovoltaic forecasting through post-processing of the global environmental multiscale numerical weather prediction model. *Progress in Photovoltaics: Research and Applications*, 21(3):284–296, 2013.
- [PKS⁺10] R. Perez, S. Kivalov, J. Schlemmer, K. Hemker Jr., D. Renné, and T. E. Hoff. Validation of short and medium term operation solar radiation forecasts in the US. *Solar Energy*, 84:2161–2172, 2010.
- [PLP⁺13] R. Perez, E. Lorenz, S. Pelland, M. Beauharnois, G. Van Knowe, K. J. Hemker, D. Heinemann, J. Remund, S. C. Muller, W. Traunmuller, G. Steinmauer, D. Pozo, J. A. Ruiz-Arias, V. Lara-Fanego, L. Ramirez-Santigosa, M. Gaston-Romero, and L. M. Pomares. Comparison of numerical weather prediction solar irradiance forecasts in the US, Canada and Europe. *Solar Energy*, 94:305–326, 2013.
- [PMW⁺07] R. Perez, K. Moore, S. Wilcox, D. Renne, and A. Zelenka. Forecasting solar radiation—preliminary evaluation of an approach based upon the National Forecast Database. *Solar Energy*, 81(6):809–812, 2007.
- [SDBF⁺15] Y. M. Saint-Drenan, S. Bolfinger, R. Fritz, S. Vogt, G. H. Good, and J. Dobschinski. An empirical approach to parameterizing photovoltaic plants for power forecasting and simulation. *Solar Energy*, 120:479–493, 2015.
- [Vap98] V. N. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, Inc., New York, 1998.
- [VMPN12] C. Voyant, M. Muselli, C. Paoli, and M.L. Nivet. Numerical weather prediction (nwp) and hybrid arma/ann model to predict global radiation. *Energy*, 39(1):341–355, 2012. Sustainable Energy and Environmental Protection 2010.

- [WH12] A. Walther and A. K. Heidinger. Implementation of the daytime cloud optical and microphysical properties algorithm (dcomp) in patmos-x. *Journal of Applied Meteorology and Climatology*, 51(7):1371–1390, 2012.
- [WKL⁺16] B. Wolff, J. Kuhnert, E. Lorenz, O. Krame, and D. Heinemann. Comparing support vector regression for PV power forecasting to a physical modeling approach using measurement, numerical weather prediction, and cloud motion data. *Solar Energy*, 135:197–208, 2016.
- [WMSZ12] F. Wang, Z. Mi, S. Su, and H. Zhao. Short-term solar irradiance forecasting model based on Artificial Neural Network using statistical feature parameters. *Energies*, 5:1355–1370, 2012.
- [ZIC14] A. Zagouras, R. H. Inman, and C. F. M. Coimbra. On the determination of coherent solar microclimates for utility planning and operators. *Solar Energy*, 102:173–188, 2014.
- [ZMAP14a] M. Zamo, O. Mestre, P. Arbogast, and O. Pannekoucke. A benchmark of statistical regression methods for short-term forecasting of photovoltaic electricity production, part I: Deterministic forecast of hourly production. *Solar Energy*, 105(0):792–803, 2014.
- [ZMAP14b] M. Zamo, O. Mestre, P. Arbogast, and O. Pannekoucke. A benchmark of statistical regression methods for short-term forecasting of photovoltaic electricity production, part II: Probabilistic forecast of daily production. *Solar Energy*, 105:804–816, 2014.
- [ZPC15] A. Zagouras, H. T. C. Pedro, and C. F. M. Coimbra. On the role of lagged exogenous variables and spatial-temporal correlations in improving the accuracy of solar forecasting methods. *Renewable Energy*, 78:203–218, 2015.